

Reflections on Tiles (in Self-Assembly)

Jacob Hendricks^{*}, Matthew J. Patitz^{**}, and Trent A. Rogers^{***}

Abstract. We define the Reflexive Tile Assembly Model (RTAM), which is obtained from the abstract Tile Assembly Model (aTAM) by allowing tiles to reflect across their horizontal and/or vertical axes. We show that the class of directed temperature-1 RTAM systems is not computationally universal, which is conjectured but unproven for the aTAM, and like the aTAM, the RTAM is computationally universal at temperature 2. We then show that at temperature 1, when starting from a single tile seed, the RTAM is capable of assembling $n \times n$ squares for n odd using only n tile types, but incapable of assembling $n \times n$ squares for n even. Moreover, we show that n is a lower bound on the number of tile types needed to assemble $n \times n$ squares for n odd in the temperature-1 RTAM. The conjectured lower bound for temperature-1 aTAM systems is $2n - 1$. Finally, we give preliminary results toward the classification of which finite connected shapes in \mathbb{Z}^2 can be assembled (strictly or weakly) by a singly seeded (i.e. seed of size 1) RTAM system, including a complete classification of which finite connected shapes be strictly assembled by a *mismatch-free* singly seeded RTAM system.

1 Introduction

Self-assembly is the process by which disorganized components autonomously combine to form organized structures. In DNA-based self-assembly, the combining ability of the components is implemented using complementary strands of DNA as the “glue”. In [22], Winfree introduced a useful mathematical model of self-assembling systems called the abstract Tile Assembly Model (aTAM) where the autonomous components are described as square tiles with specifiable glues on their edges and the attachment of these components occurs spontaneously when glues match. The aTAM provides a convenient way of describing self-assembling systems and their resulting assemblies, and serves as the underpinning of many studies of the properties of self-assembling systems. For a

^{*} Department of Computer Science and Computer Engineering, University of Arkansas, jhendric@uark.edu. Supported in part by National Science Foundation Grant CCF-1117672 and CCF-1422152.

^{**} Department of Computer Science and Computer Engineering, University of Arkansas, patitz@uark.edu. Supported in part by National Science Foundation Grant CCF-1117672 and CCF-1422152.

^{***} Department of Mathematical Sciences, University of Arkansas, tar003@email.uark.edu. This author’s research was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1450079, and National Science Foundation grants CCF-1117672 and CCF-1422152.

comprehensive survey of tile-based self-assembly including models other than the aTAM, see [5, 15].

From the broad collection of results in the aTAM, one property of systems that has been shown to yield enormous power is *cooperation*. The notion of cooperation captures the phenomenon where the attachment of a new tile to a growing assembly requires it to bind to more than one tile (usually 2) already in the assembly. The requirement for cooperation is determined by a system parameter known as the *temperature*, and when the temperature is equal to 1 (a.k.a. temperature-1 systems), there is no requirement for cooperation. A long-standing conjecture is that temperature-1 aTAM systems are in fact not capable of universal computation or efficient shape building, although it is well-known that temperature ≥ 2 systems are. However, in actual laboratory implementations of DNA-based tiles [1, 14, 19, 21, 23], the self-assembly performed by temperature-2 systems does not match the error-free behavior dictated by the aTAM, but instead, a frequent source of errors is the binding of tiles using only a single bond. Thus, temperature-1 behavior erroneously occurs and cannot be completely prevented.

Many models of self-assembly can be thought of as extensions of the aTAM (e.g. [2, 4, 6, 8, 9, 16]), and for these models it is common to study the added power that an extra property or constraint gives the extended model. For example, in [2, 6, 8, 16], it is shown that at temperature 1, when the aTAM is appropriately extended, the resulting models are computationally universal and capable of efficiently assembling shapes. In this paper, we take the opposite approach and remove a constraint that the aTAM imposes with the goal of modelling physical systems that may be incapable of enforcing these constraints. Tiles in the aTAM are not allowed to flip or rotate prior to attachment to an existing assembly. While this assumption is a realistic one for many implementations of DNA-based tiles (e.g. [22]), for certain implementations (e.g. [12]), it is unknown whether or not both of the conditions of this assumption can be physically enforced. (See [3, 10, 11, 17] for more experimentally produced building blocks and systems.) When DNA is used as the binding agent, single stranded DNA can be used to prevent relative tile rotation by encoding a direction (north/south or east/west) in the DNA sequence so that only strands with appropriately matching directions are complementary; on the other hand, preventing tiles from flipping may not always be possible, especially if the glues of different sides of a tile are encoded by disjoint DNA complexes. Therefore, we consider a model based on the aTAM where tiles may nondeterministically flip horizontally and/or vertically prior to attachment.

We introduce the *Reflexive Tile Assembly Model* (RTAM), which can be thought of as the aTAM with the relaxed constraint that tiles in the RTAM are allowed to flip horizontally and/or vertically. Also, unlike most formulations of the aTAM where complementary strands of DNA are represented with the same glue label, the RTAM explicitly uses complementary glues. This importantly prevents copies of tiles of the same type from being able to flip and bind to each other, and is the actual reality with DNA-based tiles. We then show a series of results within the RTAM. First we show that at temperature 1, the class of

directed RTAM systems – systems which yield a single pattern up to reflection and ignoring tile orientation – are only capable of assembling patterns that are essentially periodic. Then, following the thesis set forth in [7], we conclude that the temperature-1 RTAM is not computationally universal. While the inability of temperature-1 aTAM systems to compute is still only conjectured, we are able to conclusively prove it for RTAM systems, specifically by using techniques developed in [7] to study temperature-1 aTAM systems. We also show that like the aTAM at temperature 2, the class of directed temperature-2 RTAM systems is computationally universal. This shows a fundamental dividing line between the powers of RTAM temperature-1 and temperature-2 directed systems. We then turn our attention to the self-assembly of squares by singly seeded temperature-1 RTAM systems where we show that for even values of $n \in \mathbb{N}$ it is impossible to self-assemble any $n \times n$ square. This is exceptional due to the fact that it is the first demonstration of a model of tile assembly in which a finite shape is proven to be impossible to self-assemble in a directed system. Typically, any finite shape can be self-assembled by a trivial system in which a unique tile type is created for each point of the shape. However, due to the ability of tiles in the RTAM to flip, it is not possible for the RTAM systems to effectively constrain the reflections of tiles to produce such even squares without the possibility of tiles growing beyond the boundaries of the squares. However, for odd values of n and m any $n \times m$ rectangle can be self-assembled using only $\frac{n+m}{2}$ tile types, thus implying that for n odd, an $n \times n$ square can be self-assembled using only n tile types. In addition, we also show that for n odd, an $n \times n$ square cannot be self-assembled using less than n tile types (thus, n is the upper and lower bound for square assembly). This is in contrast to the aTAM at temperature 1, where the conjectured lower bound for assembling an $n \times n$ square is $2n - 1$, and hints that in certain situations the ability of RTAM tiles to attach in flipped orientations can be effectively harnessed to more efficiently build shapes than systems in the aTAM. Finally, we give preliminary results toward the classification of the finite connected shapes in \mathbb{Z}^2 that can be assembled (strictly or weakly) by a singly seeded RTAM system, including a complete classification of which finite connected shapes be strictly assembled by a *mismatch-free* singly seeded temperature-1 RTAM system. We also show that arbitrary shapes with scale factor 2 can be assembled in the singly seeded temperature-1 RTAM. These combined results show that the ability of tiles to bind in flipped orientations is sometimes provably limiting, while at other times can provide advantages, and they provide a solid framework for the study of self-assembling systems composed of molecular building blocks unable to enforce the constraints of the aTAM.

The layout of this paper is as follows. In Section 2 we present a high-level definition of the RTAM. (Due to space constraints, a more technically detailed definition can be found in Section A of the appendix, and so can the full proofs of each result.) Section 3 contains the proof that temperature-1 RTAM systems cannot perform universal computation and that temperature-2 systems can. In Section 4 we present our results related to the self-assembly of shapes in the RTAM, including our results about assembling squares and classifying the finite connected shapes that self-assemble in the RTAM.

2 Definition of the Reflexive Tile Assembly Model

The Reflexive Tile Assembly Model (RTAM) is essentially equivalent to the abstract Tile Assembly Model (aTAM) [13, 18, 20, 22] but with the modification that tiles are allowed to possibly “flip” across their horizontal and/or vertical axes before attaching to an assembly. Also, as in some formulations of the aTAM, it is assumed that glues bind to complementary versions of themselves (so that two tiles of the same type which are flipped relative to each other can’t simply bind to each other along the same but reflected side). We now give a brief definition of the RTAM. See Section A for more detailed definitions. Our notation is similar (and where appropriate, identical) to that of [13].

We work in the 2-dimensional discrete space \mathbb{Z}^2 . Define the set $U_2 = \{(0, 1), (1, 0), (0, -1), (-1, 0)\}$ to be the set of all *unit vectors* in \mathbb{Z}^2 . We also sometimes refer to these vectors by their cardinal directions N, E, S, W , respectively. All *graphs* in this paper are undirected. A *grid graph* is a graph $G = (V, E)$ in which $V \subseteq \mathbb{Z}^2$ and every edge $\{\mathbf{a}, \mathbf{b}\} \in E$ has the property that $\mathbf{a} - \mathbf{b} \in U_2$. Intuitively, a tile type t is a unit square that can be translated and flipped across its vertical and/or horizontal axes, but not rotated. This provides each tile type with a pair of North-South (NS) sides and a pair of East-West (EW) sides, such that either side $s \in NS$ may be facing north while the other is facing south (and vice versa for the EW glues). For ease of discussion, however, we will talk about tile types as being defined in fixed orientations, but then allow them to attach to assemblies in possibly flipped orientations.

Fix a finite set T of tile types. Each side of a tile t in T has a “glue” with “label” –a string over some fixed alphabet– and “strength” –a nonnegative integer– specified by its type t . Let $R = \{D, V, H, B\}$ be the set of permissible reflections for a tile which is assumed to begin in the default orientation, where D corresponds to no change from the default, V a single vertical flip (i.e. a reflection across the x -axis), H a single horizontal flip (i.e. a reflection across the y -axis), and B a single horizontal flip and a single vertical flip. See Figure 1 for an example of each. It is important to note that a glue does not have any particular orientation along the edge on which it resides, and so remains unchanged throughout reflections.




Fig. 1: Left to right: (1) Default orientation of an example tile type t , (2) t flipped vertically, (3) t flipped horizontally, (4) t flipped across both axes

Glues on adjacent edges of two tiles may bind iff they have complementary labels and the same strength. An *assembly* is a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T \times R$ defined on at least one input, with points $\mathbf{x} \in \mathbb{Z}^2$ at which $\alpha(\mathbf{x})$ is undefined interpreted to be empty space, so that α is the set of points with *oriented* tiles.

Two assemblies α and β are equivalent iff one of them can be flipped and translated so that they perfectly match at all locations. Each assembly α induces a *binding graph*, a grid graph whose vertices are positions occupied by tiles, according to α , with an edge between two vertices if the tiles at those vertices have complementary glues of equal strength. Then, for some $\tau \in \mathbb{N}$, an assembly α is τ -*stable* if every cut of the binding graph of α has weight at least τ , where the weight of an edge is the strength of the glue it represents. When τ is clear from context, we say α is *stable*. For a tile set T , we let $p_T : T \times R \rightarrow T$ be the projection map onto T (i.e. $p_T((t, r)) = t$). A *configuration* given by an assembly α is defined to be the map from \mathbb{Z}^2 to T given by $p_T \circ \alpha$.

Self-assembly begins with a *seed assembly* σ , in which each tile has a specified and fixed orientation, and proceeds asynchronously and nondeterministically, with tiles in any valid reflection in R adsorbing one at a time to the existing assembly in any manner that preserves τ -stability at all times. A *reflexive tile assembly system* (RTAS or just TAS when the context is clear) is an ordered triple $\mathcal{T} = (T, \sigma, \tau)$, where T is a finite set of tile types, σ is a seed assembly with finite domain in which each tile is given a fixed orientation, and $\tau \in \mathbb{N}$ is the *temperature* of the system which intends to model physical temperature. We use “temperature- τ system” to refer to any TAS with temperature τ . We write $\mathcal{A}[\mathcal{T}]$ for the set of all τ -stable assemblies that can arise (in finitely many steps or in the limit) in \mathcal{T} . An assembly $\alpha \in \mathcal{A}[\mathcal{T}]$ is *terminal*, and we write $\alpha \in \mathcal{A}_\square[\mathcal{T}]$, if no tile can be τ -stably added to it. It is clear that $\mathcal{A}_\square[\mathcal{T}] \subseteq \mathcal{A}[\mathcal{T}]$. We say that \mathcal{T} is *directed* if and only if all of the assemblies (upto reflection and translation) of a directed system give the same configuration.

A set $X \subseteq \mathbb{Z}^2$ *weakly self-assembles* if there exists a TAS $\mathcal{T} = (T, \sigma, \tau)$ and a set $B \subseteq T$ such that for each $\alpha \in \mathcal{A}_\square[\mathcal{T}]$ there exists a reflection $r \in R$ and a translation $\mathbf{v} \in \mathbb{Z}^2$ such that for the assembly, α^r say, corresponding to α reflected according to r then translated by \mathbf{v} and $\alpha_r^{-1}(B) = X$ holds. Essentially, weak self-assembly can be thought of as the creation (or “painting”) of a pattern of tiles from B (usually taken to be a unique “color” such as black) on a possibly larger “canvas” of un-colored tiles. Also, a set X *strictly self-assembles* if there is a TAS \mathcal{T} such that for each assembly $\alpha \in \mathcal{A}_\square[\mathcal{T}]$ there exists a reflection $r \in R$ and a translation $\mathbf{v} \in \mathbb{Z}^2$ such that $\alpha_r = F(\alpha, r, \mathbf{v})$ and $\text{dom } \alpha_r = X$. Essentially, strict self-assembly means that tiles are only placed in positions defined by the shape. Note that if X strictly self-assembles, then X weakly self-assembles.

In this paper, we also consider scaled-up versions of shapes. Formally, if X is a shape and $c \in \mathbb{N}$, then a *c-scaling* of X is defined as the set $X^c = \{(x, y) \in \mathbb{Z}^2 \mid (\lfloor \frac{x}{c} \rfloor, \lfloor \frac{y}{c} \rfloor) \in X\}$. Intuitively, X^c is the shape obtained by replacing each point in X with a $c \times c$ block of points. We refer to the natural number c as the *scale factor*.

3 The RTAM is Not Computationally Universal at $\tau = 1$

In this section, we show that directed RTAM systems are not computationally universal by showing that any shape weakly assembled by a directed RTAM

system is “simple”. We will first define our notion of simple. Many of the following definitions can also be found in [7].

Definition 1. A set $X \subseteq \mathbb{Z}^2$ is semi-doubly periodic if there exist three vectors \mathbf{b} , \mathbf{u} , and \mathbf{v} in \mathbb{Z}^2 such that

$$X = \{ \mathbf{b} + n \cdot \mathbf{u} + m \cdot \mathbf{v} \mid n, m \in \mathbb{N} \}.$$

Less formally, a semi-doubly periodic set is a set that repeats infinitely along two vectors (linearly independent vectors in the non-degenerate case), starting at some base point \mathbf{b} . Now, let $\mathcal{T} = (T, \sigma, 1)$ refer to a directed, temperature-1 RTAM system. We show that any such \mathcal{T} weakly self-assembles a set $X \subseteq \mathbb{Z}^2$ that is a finite union of semi-doubly periodic sets.

Theorem 1. Let $\mathcal{T} = (T, \sigma, 1)$ be a directed RTAM system. If a set $X \subseteq \mathbb{Z}^2$ weakly self-assembles in \mathcal{T} , then X is a finite union of semi-doubly periodic sets.

Proof. (sketch) Here we give a high-level sketch of the proof of Theorem 1. See Section B for a rigorous proof. The basic idea of the proof is as follows. For an RTAM system $\mathcal{T} = (T, \sigma, 1)$ we consider all of the paths of n tiles (for n to be defined) that can assemble from each exposed glue of σ such that each consecutive tile that binds forming the path attach via a north or west glue (and we say that such a path “extends to the north-west”).




Fig. 2: A depiction of tiles for a path that can assemble in \mathcal{T} . The original path is on the left. The path on the right is a modification to the path on the left that must also be able to assemble in \mathcal{T} . The blue tiles labeled \mathbf{v}_0 and \mathbf{v}_1 are of the same tile type and orientation. The tiles labeled 1 through 6 can be repeated indefinitely as depicted in Table 1(a).

Any finite path is trivially the union of semi-doubly periodic sets. Then, for n sufficiently large, a path of n tiles that extends to the north-west must contain two distinct tiles t_1 and t_2 of the same tile type in the same orientation. If for every such path, every two distinct tiles t_1 and t_2 of the same tile type in

the same orientation lie on a horizontal or vertical line, then it can be argued that the terminal configuration of \mathcal{T} must consist of finitely many infinitely long horizontal or vertical paths connected to σ , and is therefore the finite union of semi-doubly periodic sets. On the other hand, if there is a path such that the two distinct tiles t_1 and t_2 of the same tile type in the same orientation do not lie on a horizontal or vertical path, then we argue that the terminal assembly of \mathcal{T} is the finite union of semi-doubly periodic sets as follows. First, we note that for such a path, π say, the tiles between t_1 and t_2 can be repeated indefinitely. This is shown in Table 1(a). Then we show how to modify π by reflecting tiles




Table 1: Each figure in this table depicts a possible path that can assemble in \mathcal{T} . The yellow tiles make up the seed and the green tiles are a path leading to the repeated tile that allows the path to repeat.

to obtain an infinite family of paths. Examples of such modified paths are shown in Table 1(b)-(h). Now, if a tile t belongs to one of these paths and has location l say, then, since \mathcal{T} is directed the terminal assembly of \mathcal{T} must contain a tile of the same type as t at each such location l . Finally, we note that all of these paths




Fig. 3: (a) A configuration that can be thought of as the “union” of the type-consistent assemblies depicted in Table 1. (b) A configuration of tiles that must weakly self-assemble in \mathcal{T} .

taken together form a semi-doubly periodic set. This is depicted in Figure 3a. Continuing this line of reasoning, we show that the terminal assembly of \mathcal{T} is the finite union of semi-doubly periodic sets. A portion of such an assembly is shown in Figure 3b.

□

An intuitive reason that Theorem 1 supports the conclusion that the RTAM is not computationally universal is as follows. Let $H = \{(i, x) \mid \text{program } i \text{ halts when run on input } x\}$ be the halting set and let M be a Turing machine that outputs a 1 if $(i, x) \in H$. The typical way of expressing the computation of M in tile assembly is as follows. For a fixed tileset T , a seed assembly σ encodes an “input” to the computation, while the “output” of 1 by M corresponds to the translation of some configuration being contained in the terminal assembly α_σ of $(T, \sigma, 1)$. For this sense of computation, the following corollary says that the set of seed assemblies that “output” a 1 is a recursive set (not just a recursively enumerable set). This would contradict the fact that the halting set is not recursive. This is stated in the following corollary. For a more formal statement of this corollary and a proof, see Section B.

Corollary 1. *For any tileset T in the RTAM and fixed finite configuration C , let S be the set of seed assemblies σ such that (1) the RTAM system $(T, \sigma, 1)$ is directed and (2) the terminal assembly of \mathcal{T} contains C . Then, S is a recursive set.*

3.1 Universal computation at $\tau = 2$

In this section we give a theorem that states that universal computation is possible in the RTAM at temperature 2. We give an example of simulating a binary counter in the RTAM and give the general proof in Section C in the appendix.

Theorem 2. *The RTAM is computationally universal at $\tau = 2$. Moreover, the class of directed RTAM systems is computationally universal at $\tau = 2$.*

First we note that given a Turing machine M , we use Lemma 7 of [2] to obtain a tile set which simulates M using a *zig-zag system*. In fact, as noted in [16], we can find a singly seeded *compact zig-zag system* $\mathcal{T} = (T, \sigma, 2)$ with $\mathcal{A}_\square[\mathcal{T}] = \{\alpha\}$ which simulates M . Then the proof of Theorem 2 relies on showing that any compact zig-zag system in the aTAM at temperature 2 can be converted into a directed RTAM system \mathcal{S} that is “almost” compact zig-zag. The RTAM system that we construct differs from a compact zig-zag system in that when the length of a row of the growing zig-zag assembly increases by a tile, a strength-2 glue is exposed that allows a tile to bind below the row. This results in the possibility of a single “misplaced” tile per row, but nevertheless, this is enough to simulate a Turing machine. The proof of Theorem 2 can be found in Section C

4 Self-assembly of shapes in the RTAM at $\tau = 1$

In this section, we discuss the self-assembly of shapes in the RTAM, especially the commonly used benchmark of squares. At temperature 2, using a zig-zag binary

counter similar to that used in Section 3.1, $n \times n$ squares can be built using the optimal $\log n / (\log \log n)$ tile types following the construction of [20] with only trivial modifications. Similarly, the majority of shapes which can be weakly self-assembled in the temperature-2 aTAM can be built in the temperature-2 RTAM, although shapes with single-tile-wide branches which are not symmetric are impossible to strictly self-assemble in the RTAM.

At temperature-1, however, the differences between the powers of the aTAM and RTAM appear to increase. Here we will demonstrate that squares whose sides are of even length cannot weakly (or therefore strictly) self-assemble in the RTAM at $\tau = 1$, although any square can strictly self-assemble in the $\tau = 1$ aTAM. We then prove a tight bound of n tile types required to self-assemble an $n \times n$ square for odd n in the RTAM at $\tau = 1$. (Which is, interestingly, better than the conjectured lower bound of $2n - 1$ for the $\tau = 1$ aTAM.)

4.1 For even n , no $n \times n$ square self-assembles in the $\tau = 1$ RTAM

Theorem 3. *For all $n \in \mathbb{Z}^+$ where n is even, there exists no RTAM system $\mathcal{T} = (T, \sigma, 1)$ where $|\sigma| = 1$ and \mathcal{T} weakly (or strictly) self-assembles an $n \times n$ square.*

Proof. (sketch) We prove Theorem 3 by contradiction, and here give a sketch of the proof. (See Section D for the full proof.) Therefore, assume that for some $n \in \mathbb{Z}^+$ such that $(n \bmod 2) = 0$, there exists an RTAM system $\mathcal{T} = (T, \sigma, 1)$ such that $|\sigma| = 1$ and \mathcal{T} weakly self-assembles an $n \times n$ square S . We take $\alpha \in \mathcal{A}_{\square}[\mathcal{T}]$ and consider the corners of the square which is weakly self-assembled.




Fig. 4: (a) Example paths in an 8×8 square (i.e. of even dimension). The path from corners **a** and **c** is composed of tiles of different colors. The path from corner **d** to a point on that path is dark grey. The path from the seed to that intersection is in black. The path to be stretched out is outlined in red. (b) The stretched out version of the paths.

There must exist a path which connects two diagonal corners, and that path must travel through 3 quadrants of the square since the dimensions are even

length and diagonal paths are not possible in the grid graph of an assembly. We then find a path connecting the corner of the quadrant (possibly) unvisited by that path, and note that since α must be connected that we can find a new path which connects the new corner to one of the original corners via a path which crosses across the midpoint of the square from either corner on the new path. (See Figure 4 (a) for an example.) Finally, we demonstrate that there is an assembly sequence which starts from the seed and grows to that new path, then builds that path in a way such that it is maximally “stretched” out by appropriately flipping the tiles. (See Figure 4 (b) for an example.) This stretched out path is producible by a valid assembly sequence but must grow beyond the bounds of the square (by at least one position), so \mathcal{T} does not weakly (or strictly) self-assemble the square. \square

4.2 Tight bounds on the tile complexity of squares of odd dimension

In this section, we prove tight bounds on the number of tiles necessary to self-assemble a square of odd dimension.

Theorem 4. *For all $n \in \mathbb{Z}^+$ where n is odd, an $n \times n$ square strictly self-assembles in an RTAM system $\mathcal{T} = (T, \sigma, 1)$ where $|T| = n$ and $|\sigma| = 1$.*

We prove Theorem 4 by giving a scheme for obtaining the tileset for any given n that exploits the fact that for n odd, an $n \times n$ square in \mathbb{Z}^2 is symmetric across a row and column points. Although Theorem 4 pertains to squares, a simple modification of the proof shows the following corollary.

Corollary 2. *For all $n, m \in \mathbb{Z}^+$ where n is odd, an $n \times m$ square strictly self-assembles in an RTAM system $\mathcal{T} = (T, \sigma, 1)$ where $|T| = \frac{n+m}{2}$ and $|\sigma| = 1$.*

We also prove that the upper bound of Theorem 4 is tight, i.e. an $n \times n$ square, where n is odd, cannot be self-assembled using less than n tile types. The proof of this theorem can be found in Section F.

Theorem 5. *For all $n \in \mathbb{Z}^+$ where n is odd, there exists no RTAM system $\mathcal{T} = (T, \sigma, 1)$ where $|T| < n$ and $|\sigma| = 1$ such that \mathcal{T} weakly (or strictly) self-assembles an $n \times n$ square.*

4.3 Assembling Finite Shapes in the RTAM

In this section we first give a corollary of Theorem 4 showing that sufficiently symmetric shapes weakly self-assemble in the RTAM. Then we prove 3 theorems about assembling finite shapes in \mathbb{Z}^2 in the RTAM. These theorems show that the assembly of finite shapes in singly seeded RTAM systems is quite a bit different than the assembly of finite shapes in the aTAM by singly seeded systems.

Given a shape S in \mathbb{Z}^2 , let χ_S denote the characteristic function of the set S . That is, $\chi_S(x, y) = 1$ if $x \in S$ and $\chi_S(x, y) = 0$ otherwise. Then, we say that a shape S is *odd-symmetric* with respect to a horizontal line $y = l$ (respectively, vertical line $x = l$) for $l \in \mathbb{Z}^2$ iff for all $(a, b) \in \mathbb{Z}^2$, $\chi_S(a, l - b) = \chi_S(a, l + b)$ (respectively, $\chi_S(l - a, b) = \chi_S(l + a, b)$). If there exists a line such that a shape,

S , is odd-symmetric with respect to this line, we say that the shape is odd-symmetric. Given a shape S , we call the smallest rectangle of points in \mathbb{Z}^2 containing S the *bounding-box* for S .

Let R denote the bounding box of an odd-symmetric shape S , and let n and m in \mathbb{N} be the dimensions of R . A simple modification to the proof of Theorem 4 where a tile is labeled with a label B if and only if it corresponds to points in S , shows that odd-symmetric shapes can weakly assemble in the RTAM.

Corollary 3. *Given a shape S in \mathbb{Z}^2 , if S is odd-symmetric, then there exists an RTAS $\mathcal{T} = (T, \sigma, \tau)$ such that $|\sigma| = 1$, $\tau \geq 1$, and \mathcal{T} weakly assembles S .*

Additionally, if one is willing to build 2 mirrored copies of the shape in each assembly, then any finite shape can be weakly self-assembled in the RTAM at $\tau = 1$, along with its mirrored copy (at a cost of tile complexity approximately equal to the number of points in the shape) by simply building a central column (or row) from which identical copies of hardcoded rows (or columns) grow, so that each side grows a reflected copy of the shape in hardcoded slices.

We say that a TAS (in either the aTAM or the RTAM) \mathcal{T} is called *mismatch-free* if for every producible assembly $\alpha \in \mathcal{A}[\mathcal{T}]$ with two neighboring tiles with abutting edges e_1 and e_2 , either e_1 and e_2 do not have glues or e_1 and e_2 have glues with matching labels and strengths. Then, for singly seeded aTAM systems, any finite connected shape can be strictly assembled by a mismatch-free system. Theorems 6, 7, and 8 show that assembling shapes in the RTAM is more complex. The proofs of these theorems can be found in Section G.

Theorem 6. *There exists a finite connected shape S in \mathbb{Z}^2 that weakly self-assembles in a singly seeded RTAM system such that there exists no singly seeded RTAM system that strictly self-assembles S .*

Theorem 7. *There exists a finite shape S in \mathbb{Z}^2 that can be strictly self-assembled by some singly seeded RTAM system such that every singly seeded RTAM system at temperature 1 which strictly self-assembles S is not directed.*

Theorem 8. *There exists a finite shape S in \mathbb{Z}^2 such that every singly seeded RTAM system that strictly self-assembles S is not mismatch-free.*

4.4 Mismatch-free Assembly of Finite Shapes in the RTAM

Given a shape S , i.e. a finite connected subset of \mathbb{Z}^2 , we say that a *graph of S* is a graph $G_S = (V, E)$ with a vertex at the center of each point in S and an edge between every pair of vertices at adjacent points of S . A *tree of S* , T_S , is a graph of S which is a tree. (See Figure 5 for examples of S , G_S , and T_S .) Given a graph $G = (V, E)$, we say that an *axis* of G is a horizontal or vertical line of vertices such that there is an edge between each pair of adjacent points on that line. Notice that two distinct axes can be collinear. Given an axis a , an *axial branch* of T_S is a branch of T_S which contains exactly one vertex v on a and all vertices and edges of T_S which are connected to a vertex that does not lie on a and is adjacent to v . We say that the branch *begins from* v . Intuitively, an axial branch is a connected component extending from an axis. (See the pink




Fig. 5: An example ϵ -symmetric shape.

highlighted portion of Figure 5c for an example axial branch off of the axis shown in green.)

A tree T_S is symmetric across an axis a if, for every vertex v contained on a , the branches of a which begin from v are symmetric across a . A tree T_S is *off-by-one symmetric* across an axis a if, for every vertex v except for at most 1, the branches of a which begin from v are symmetric across a . See Figure 5c for an example of such a tree, with the axis a shown in green.

Definition 2. A tree T is ϵ -symmetric if and only if for any axis a of T , T is off-by-one symmetric across a .

Definition 3. Given a shape S with graph G_S , we say that S is ϵ -symmetric if and only if there exists a spanning tree, T_S , of G_S such that T_S is ϵ -symmetric.

For an example of an ϵ -symmetric shape S , see Figure 5a. The tree T_S is off-by-one symmetric across the vertical green axis, the branches off of that axis are symmetric across the horizontal yellow axes, and the branches off of those axes are symmetric across the vertical blue axes. The following theorem gives a complete classification of finite connected shapes which can be assembled by temperature-1 singly seeded mismatch-free RTAM systems. The proof of this theorem is given in Section H.

Theorem 9. Let $S \subset \mathbb{Z}^2$ be a finite connected shape. There exists a mismatch-free RTAM system $\mathcal{T} = (T, \sigma, 1)$ with $|\sigma| = 1$ that strictly assembles S if and only if S is ϵ -symmetric.

While Theorem 9 shows exactly which shapes can be assembled without cooperation or mismatches by singly seeded RTAM systems, the following theorem shows that with cooperation, RTAM systems can assemble arbitrary scale factor 2 shapes. The proof of this theorem is in Section I.

Theorem 10. Let $S \subset \mathbb{Z}^2$ be a finite connected shape, and S^2 be S at scale factor 2. There exists a mismatch-free RTAM system $\mathcal{T} = (T, \sigma, 2)$ with $|\sigma| = 1$ that strictly self-assembles S^2 .

References

1. Robert D. Barish, Rebecca Schulman, Paul W. K. Rothmund, and Erik Winfree. An information-bearing seed for nucleating algorithmic self-assembly. *Proceedings of the National Academy of Sciences*, 106(15):6054–6059, April 2009.

2. Matthew Cook, Yunhui Fu, and Robert T. Schweller. Temperature 1 self-assembly: Deterministic assembly in 3D and probabilistic assembly in 2D. In *SODA 2011: Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2011.
3. Cristina Costa Santini, Jonathan Bath, Andy M. Tyrrell, and Andrew J. Turberfield. A clocked finite state machine built from dna. *Chem. Commun.*, 49:237–239, 2013.
4. E. D. Demaine, M. L. Demaine, S. P. Fekete, M. J. Patitz, R. T. Schweller, A. Winslow, and D. Woods. One tile to rule them all: Simulating any tile assembly system with a single universal tile. In J. Esparza, P. Fraigniaud, T. Husfeldt, and E. Koutsoupias, editors, *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP 2014)*, IT University of Copenhagen, Denmark, July 8–11, 2014, volume 8572 of *LNCS*, pages 368–379. Springer Berlin Heidelberg, 2014.
5. David Doty. Theory of algorithmic self-assembly. *Commun. ACM*, 55(12):78–88, December 2012.
6. David Doty, Lila Kari, and Benoît Masson. Negative interactions in irreversible self-assembly. *Algorithmica*, 66(1):153–172, 2013.
7. David Doty, Matthew J. Patitz, and Scott M. Summers. Limitations of self-assembly at temperature 1. *Theoretical Computer Science*, 412:145–158, 2011.
8. Sándor P. Fekete, Jacob Hendricks, Matthew J. Patitz, Trent A. Rogers, and Robert T. Schweller. Universal computation with arbitrary polyomino tiles in non-cooperative self-assembly. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2015)*, San Diego, CA, USA January 4–6, 2015, pages 148–167.
9. Bin Fu, Matthew J. Patitz, Robert T. Schweller, and Robert Sheline. Self-assembly with geometric tiles. In *Proceedings of the 39th International Colloquium on Automata, Languages and Programming*, ICALP, pages 714–725, 2012.
10. Dongran Han, Suchetan Pal, Yang Yang, Shuoxing Jiang, Jeanette Nangreave, Yan Liu, and Hao Yan. Dna gridiron nanostructures based on four-arm junctions. *Science*, 339(6126):1412–1415, 2013.
11. Yonggang Ke, Luvena L Ong, William M Shih, and Peng Yin. Three-dimensional structures self-assembled from dna bricks. *Science*, 338(6111):1177–1183, 2012.
12. Jin-Woo Kim, Jeong-Hwan Kim, and Russell Deaton. Dna-linked nanoparticle building blocks for programmable matter. *Angewandte Chemie International Edition*, 50(39):9185–9190, 2011.
13. James I. Lathrop, Jack H. Lutz, and Scott M. Summers. Strict self-assembly of discrete Sierpinski triangles. *Theoretical Computer Science*, 410:384–405, 2009.
14. Chengde Mao, Thomas H. LaBean, John H. Relf, and Nadrian C. Seeman. Logical computation using algorithmic self-assembly of DNA triple-crossover molecules. *Nature*, 407(6803):493–6, 2000.
15. Matthew J. Patitz. An introduction to tile-based self-assembly and a survey of recent results. *Natural Computing*, 13(2):195–224, 2014.
16. Matthew J. Patitz, Robert T. Schweller, and Scott M. Summers. Exact shapes and tiling universality at temperature 1 with a single negative glue. In *Proceedings of the 17th international conference on DNA computing and molecular programming*, DNA’11, pages 175–189, Berlin, Heidelberg, 2011. Springer-Verlag.
17. Andre V Pinheiro, Dongran Han, William M Shih, and Hao Yan. Challenges and opportunities for structural dna nanotechnology. *Nature nanotechnology*, 6(12):763–772, 2011.
18. Paul W. K. Rothemund. *Theory and Experiments in Algorithmic Self-Assembly*. PhD thesis, University of Southern California, December 2001.

19. Paul W. K. Rothemund, Nick Papadakis, and Erik Winfree. Algorithmic self-assembly of dna sierpinski triangles. *PLoS Biol*, 2(12):e424, 12 2004.
20. Paul W. K. Rothemund and Erik Winfree. The program-size complexity of self-assembled squares (extended abstract). In *STOC '00: Proceedings of the thirty-second annual ACM Symposium on Theory of Computing*, pages 459–468, Portland, Oregon, United States, 2000. ACM.
21. Rebecca Schulman and Erik Winfree. Synthesis of crystals with a programmable kinetic barrier to nucleation. *Proceedings of the National Academy of Sciences*, 104(39):15236–15241, 2007.
22. Erik Winfree. *Algorithmic Self-Assembly of DNA*. PhD thesis, California Institute of Technology, June 1998.
23. Erik Winfree, Furong Liu, Lisa A. Wenzler, and Nadrian C. Seeman. Design and self-assembly of two-dimensional DNA crystals. *Nature*, 394(6693):539–44, 1998.

A Formal Definition of the Reflexive Tile Assembly Model

We work in the 2-dimensional discrete space \mathbb{Z}^2 . Define the set $U_2 = \{(0,1), (1,0), (0,-1), (-1,0)\}$ to be the set of all *unit vectors* in \mathbb{Z}^2 . We also sometimes refer to these vectors by their cardinal directions N, E, S, W , respectively. All *graphs* in this paper are undirected. A *grid graph* is a graph $G = (V, E)$ in which $V \subseteq \mathbb{Z}^2$ and every edge $\{\mathbf{a}, \mathbf{b}\} \in E$ has the property that $\mathbf{a} - \mathbf{b} \in U_2$.

Intuitively, a tile type t is a unit square that can be translated and flipped across its vertical and/or horizontal axes, but not rotated. This provides each tile type with a pair of North-South (NS) sides and a pair of East-West (EW) sides, such that either side $s \in NS$ may be facing north while the other is facing south (and vice versa for the EW glues). For ease of discussion, however, we will talk about tile types as being defined in fixed orientations, but then allow them to attach to assemblies in possibly flipped orientations. Therefore, we define each t as having a well-defined “side \mathbf{u} ” for each $\mathbf{u} \in U_2$. Each side \mathbf{u} of t has a “glue” with “label” $\text{label}_t(\mathbf{u})$ —a string over some fixed alphabet—and “strength” $\text{str}_t(\mathbf{u})$ —a nonnegative integer—specified by its type t . Let $R = \{D, V, H, B\}$ be the set of permissible reflections for a tile which is assumed to begin in the default orientation, where D corresponds to no change from the default, V a single vertical flip (i.e. a reflection across the x -axis), H a single horizontal flip (i.e. a reflection across the y -axis), and B a single horizontal flip and a single vertical flip. (Note that the ordering of flips for B does not matter as either ordering results in the same orientation, and also that all combinations of possibly many flips across each axis result only in tiles of the orientations provided by R .) See Figure 6 for an example of each. Let $S : R \times U_2 \rightarrow U_2$ be a function which takes a type of reflection $r \in R$ and a side $s \in U_2$, and which returns the side of a tile in its default orientation which would appear on side s of the tile when it has been reflected according to r . (E.g. for the tile type shown in Figure 6, $S(H, W) = E$, and $S(H, N) = N$.) It is important to note that a glue does not have any particular orientation along the edge on which it resides, and so remains unchanged throughout reflections.




Fig. 6: Left to right: (1) Default orientation of an example tile type t , (2) t flipped vertically, (3) t flipped horizontally, (4) t flipped across both axes

Two tiles t and t' that are placed at the points \mathbf{a} and $\mathbf{a} + \mathbf{u}$ and reflected by $r \in R$ and $r' \in R$, respectively, *bind* with *strength* $\text{str}_t(S(r, \mathbf{u}))$ if and only if $(\text{label}_t(S(r, \mathbf{u})), \text{str}_t(S(r, \mathbf{u}))) = (\overline{\text{label}_{t'}(S(r', -\mathbf{u}))}, \text{str}_{t'}(S(r', -\mathbf{u})))$. That

is, the glues on adjacent edges of two tiles bind iff they have complementary labels and the same strength.

In the subsequent definitions, given two partial functions f, g , we write $f(x) = g(x)$ if f and g are both defined and equal on x , or if f and g are both undefined on x .

Fix a finite set T of tile types. A T -assembly, sometimes denoted simply as an assembly when T is clear from the context, is a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T \times R$ defined on at least one input, with points $\mathbf{x} \in \mathbb{Z}^2$ at which $\alpha(\mathbf{x})$ is undefined interpreted to be empty space, so that $\text{dom } \alpha$ is the set of points with *oriented* tiles. We write $|\alpha|$ to denote $|\text{dom } \alpha|$, and we say α is *finite* if $|\alpha|$ is finite. For a given location $\mathbf{v} \in \mathbb{Z}^2$, we denote the tile in α at location \mathbf{v} by $\alpha(\mathbf{v})$ (if no tile exists there, $\alpha(\mathbf{v})$ is undefined). Let F be a function which takes as input an assembly α , a reflection $r \in R$, and a translation vector $\mathbf{v} \in \mathbb{Z}^2$, and which returns the assembly α^r , corresponding to α reflected according to r then translated by \mathbf{v} . We say that two assemblies, α and β , are equivalent iff there exists some reflection $r \in R$ and translation vector $\mathbf{v} \in \mathbb{Z}^2$ such that for $\beta' = F(\beta, r, \mathbf{v})$, $|\alpha| = |\beta'|$ and for all $\mathbf{v} \in |\alpha|$, $\alpha(\mathbf{v}) = \beta'(\mathbf{v})$. That is, α and β are equivalent iff one of them can be flipped and translated so that they perfectly match at all locations. For assemblies α and α' , we say that α is a *subassembly* of α' , and write $\alpha \sqsubseteq \alpha'$, if $\text{dom } \alpha \subseteq \text{dom } \alpha'$ and $\alpha(\mathbf{x}) = \alpha'(\mathbf{x})$ for all $\mathbf{x} \in \text{dom } \alpha$. An assembly α is τ -stable For some $\tau \in \mathbb{N}$, an assembly α is τ -stable if every cut of the binding graph of α has weight at least τ , where the weight of an edge is the strength of the glue it represents. When τ is clear from context, we say α is *stable*.

For a tile set T , we let $p_T : T \times R \rightarrow T$ be the projection map onto T (i.e. $p_T((t, r)) = t$). A *configuration* given by an assembly α is defined to be the map from \mathbb{Z}^2 to T given by $p_T \circ \alpha$.

Self-assembly begins with a *seed assembly* σ , in which each tile has a specified and fixed orientation, and proceeds asynchronously and nondeterministically, with tiles in any valid reflection in R adsorbing one at a time to the existing assembly in any manner that preserves τ -stability at all times. A *tile assembly system* (TAS) is an ordered triple $\mathcal{T} = (T, \sigma, \tau)$, where T is a finite set of tile types, σ is a seed assembly with finite domain in which each tile is given a fixed orientation, and $\tau \in \mathbb{N}$ is the *temperature*. A *generalized tile assembly system* (GTAS) is defined similarly, but without the finiteness requirements. We write $\mathcal{A}[\mathcal{T}]$ for the set of all assemblies that can arise (in finitely many steps or in the limit) from \mathcal{T} . An assembly $\alpha \in \mathcal{A}[\mathcal{T}]$ is *terminal*, and we write $\alpha \in \mathcal{A}_\square[\mathcal{T}]$, if no tile can be τ -stably added to it. It is clear that $\mathcal{A}_\square[\mathcal{T}] \subseteq \mathcal{A}[\mathcal{T}]$.

An assembly sequence in a TAS \mathcal{T} is a (finite or infinite) sequence $\alpha = (\alpha_0, \alpha_1, \dots)$ of assemblies in which each α_{i+1} is obtained from α_i by the addition of a single tile. The *result* $\text{res}(\alpha)$ of such an assembly sequence is its unique limiting assembly. (This is the last assembly in the sequence if the sequence is

finite.) The set $\mathcal{A}[\mathcal{T}]$ is partially ordered by the relation \longrightarrow defined by

$$\alpha \longrightarrow \alpha' \text{ iff there is an assembly sequence } \boldsymbol{\alpha} = (\alpha_0, \alpha_1, \dots) \\ \text{such that } \alpha_0 = \alpha \text{ and } \alpha' = \text{res}(\boldsymbol{\alpha}).$$

We say that \mathcal{T} is *strongly directed* if and only if either $|\mathcal{A}_\square[\mathcal{T}]| = 1$ or if for every pair of terminal assemblies $\alpha, \beta \in \mathcal{A}_\square[\mathcal{T}]$, there exists a reflection $r \in R$ and a translation vector $\mathbf{v} \in \mathbb{Z}^2$ such that $\alpha = F(\beta, r, \mathbf{v})$. Furthermore, we say that \mathcal{T} is *directed* if and only if for all $\alpha, \beta \in \mathcal{T}$, there exists a reflection $r \in R$ and a translation vector $\mathbf{v} \in \mathbb{Z}^2$ such that $p_T \circ \alpha = p_T \circ F(\beta, r, \mathbf{v})$. In other words, all of the assemblies of a directed systems give the same configuration.

A set $X \subseteq \mathbb{Z}^2$ *weakly self-assembles* if there exists a TAS $\mathcal{T} = (T, \sigma, \tau)$ and a set $B \subseteq T$ such that for each $\alpha \in \mathcal{A}_\square[\mathcal{T}]$ there exists a reflection $r \in R$ and a translation $\mathbf{v} \in \mathbb{Z}^2$ such that $\alpha_r = F(\alpha, r, \mathbf{v})$ and $\alpha_r^{-1}(B) = X$ holds. Essentially, weak self-assembly can be thought of as the creation (or “painting”) of a pattern of tiles from B (usually taken to be a unique “color” such as black) on a possibly larger “canvas” of un-colored tiles.

A set X *strictly self-assembles* if there is a TAS \mathcal{T} such that for each assembly $\alpha \in \mathcal{A}_\square[\mathcal{T}]$ there exists a reflection $r \in R$ and a translation $\mathbf{v} \in \mathbb{Z}^2$ such that $\alpha_r = F(\alpha, r, \mathbf{v})$ and $\text{dom } \alpha_r = X$. Essentially, strict self-assembly means that tiles are only placed in positions defined by X . Note that if X strictly self-assembles, then X weakly self-assembles. X in the definition of strict or weak self-assembly is called a *shape* in \mathbb{Z}^2 .

In this paper, we also consider scaled-up versions of shapes. Formally, if X is a shape and $c \in \mathbb{N}$, then a *c-scaling* of X is defined as the set $X^c = \{(x, y) \in \mathbb{Z}^2 \mid (\lfloor \frac{x}{c} \rfloor, \lfloor \frac{y}{c} \rfloor) \in X\}$. Intuitively, X^c is the shape obtained by replacing each point in X with a $c \times c$ block of points. We refer to the natural number c as the *scale factor*.

A.1 Paths in the Binding Graph and as Assemblies

Given an assembly α and locations \mathbf{x} and \mathbf{y} such that $\mathbf{x}, \mathbf{y} \in \text{dom } \alpha$, we define a *path in α from \mathbf{x} to \mathbf{y}* (or simply a *path from \mathbf{x} to \mathbf{y}*) as a simple directed path in the binding graph of α with the first location being \mathbf{x} and the last \mathbf{y} . We refer to such a path as π_x^y , and for $k = |\pi_x^y|$ (i.e. k is the length of, or number of tiles on, π_x^y) and $0 \leq i < k$, let $\pi_x^y(i)$ be the i th location of π_x^y . Thus, $\pi_x^y(0) = \mathbf{x}$, and $\pi_x^y(k-1) = \mathbf{y}$. We can thus refer to the i th tile on π_x^y and its reflection as $\alpha(\pi_x^y(i))$, and as shorthand will often refer to locations and/or tiles along a path. Regardless of the order in which the tiles of π_x^y were placed in α , we define *input* and *output* sides for each tile in π_x^y (except for the first and last, respectively) in relation to their position on π_x^y . The input side of the i th tile of π_x^y , $\alpha(\pi_x^y(i))$, is that which binds to $\alpha(\pi_x^y(i-1))$, and the output side is that which binds to $\alpha(\pi_x^y(i+1))$. We denote these sides as $IN(\pi_x^y(i))$ and $OUT(\pi_x^y(i))$, respectively. (Thus, $\alpha(\pi_x^y(0))$ has no input side, and $\alpha(\pi_x^y(k-1))$ has no output side.) Note that in a temperature-1 system, an assembly α' exactly representing π_x^y would

be able to grow solely from $\alpha(\pi_x^y(0))$, in the order of π_x^y , with each tile having input and output sides as defined for π_x^y .

B Proof of Theorem 1 and Additional Corollaries

In the proof of Theorem 1, we will be considering many different assemblies and will use the following definition to form a “union” of the configurations given by these assemblies.

Definition 4. *Two assemblies α and β are type-consistent iff for each $\mathbf{l} \in \text{dom } \alpha \cap \text{dom } \beta$, $p_T(\alpha(\mathbf{l})) = p_T(\beta(\mathbf{l}))$.*

Recall that p_T is the projection from $T \times R$ onto T . Notice that if two assemblies α and β are type-consistent, then we can give a well-defined partial function f from \mathbb{Z}^2 to T by $x \mapsto t$ iff either $x \in \text{dom } \alpha$ and $\alpha(x) = t$ or $x \in \text{dom } \beta$ and $\beta(x) = t$, and otherwise f is undefined. The following definitions will also be useful in the proof of Theorem 1.

Let $c \in \mathbb{N}$ and $\mathbf{v} \in \mathbb{Z}^2$. The *box of radius c centered about the point \mathbf{v}* is the set of points defined as $B_c(\mathbf{v}) = \{ (x, y) + \mathbf{v} \mid |x| \leq c \text{ and } |y| \leq c \}$. Finally, we say that a path in the binding graph *turns* if there are two nodes $\mathbf{n}_0 = (x_0, y_0)$ and $\mathbf{n}_1 = (x_1, y_1)$ in the path such that $x_0 \neq x_1$ and $y_0 \neq y_1$. In other words, a path turns if its vertices do not lie on a horizontal or vertical line. A path of tiles in an assembly is said to turn if the corresponding path in the binding graph turns.

Without loss of generality, assume that σ contains the point $\mathbf{0} \in \mathbb{Z}^2$ where $\mathbf{0} = (0, 0)$. Let $c \in \mathbb{N}$ be the minimal number such that $\sigma \subseteq B_c(\mathbf{0})$ and let $c' \in \mathbb{N}$ be such that $c' = 4|T| + c + 1$. Now let $\alpha \in \mathcal{A}_{\square}[T]$ be a terminal assembly such that α and σ are type-consistent. Note that a finite assembly trivially weakly assembles a finite union of semi-doubly periodic sets. Therefore, we only consider the case where $\overline{B_{c'}(\mathbf{0})} \cap \text{dom } \alpha \neq \emptyset$. Let $\mathbf{x} \in \mathbb{Z}^2$ be a tile location for some tile in $\overline{B_{c'}(\mathbf{0})} \cap \text{dom } \alpha$, and let $\pi_{\sigma}^{\mathbf{x}}$ be a path in G_{α} from a node corresponding to a tile in σ to \mathbf{x} . (Notice the abuse of notation here. We are using σ in the notation $\pi_{\sigma}^{\mathbf{x}}$ where we should be using the location of the tile in σ .) Notice that $\pi_{\sigma}^{\mathbf{x}}$ must contain at least $4|T| + 1$ nodes. Now, at temperature 1, $\sigma \cup \pi_{\sigma}^{\mathbf{x}}$ is a producible assembly of T ; denote this assembly by β . We will modify β by modifying the path $\pi_{\sigma}^{\mathbf{x}}$. Let β be the assembly sequence for β , and let t be the first tile placed outside of $B_c(\mathbf{0})$ using this assembly sequence. This tile is marked as the red tile in Figure 7. We will denote the path from t to \mathbf{x} by $\pi_t^{\mathbf{x}}$, and denote the subpath of $\pi_{\sigma}^{\mathbf{x}}$ that does not contain the tiles of $\pi_t^{\mathbf{x}}$ by π_{σ}^t . (Again we abuse notion here. We are using t in the notation $\pi_t^{\mathbf{x}}$ where we should technically use the location of t .) Without loss of generality, suppose that for $\mathbf{l} = (l_x, l_y)$, the location of t , for all $(x, y) \in B_c(\mathbf{0})$, $l_y > y$. In other words, suppose that t lies to the north of $B_c(\mathbf{0})$. This is depicted in Figure 7. Now we modify the path $\pi_t^{\mathbf{x}}$ as follows.

Let $\alpha = (\alpha_0, \alpha_1, \dots, \alpha_n)$ be the assembly sequence such that $\alpha_0 = \sigma$ and α_i is obtained from α_{i-1} by the attachment of the tile $\pi_{\sigma}^{\mathbf{x}}(i-1)$ of $\pi_{\sigma}^{\mathbf{x}}$. We modify α as follows. Suppose that k is such that t is the tile at the location in




Fig. 7: A depiction of tiles for a path that can assemble in \mathcal{T} . The original path is on the left. The path on the right is a modification to the path on the left that must also be able to assemble in \mathcal{T} . On the left, the light grey inner box delineates $B_c(\mathbf{0})$ and the outer one delineates $B_{c'}(\mathbf{0})$. The blue tiles labeled \mathbf{v}_0 and \mathbf{v}_1 are at the respective locations \mathbf{v}_0 and \mathbf{v}_1 in \mathbb{Z}^2 . The red tile is the first tile (t) along the path with location outside of $B_c(\mathbf{0})$ and the black tile is the first tile along the path with location (\mathbf{x}) outside of $B_{c'}(\mathbf{0})$.

$\text{dom } \alpha_k \setminus \text{dom } \alpha_{k-1}$. That is, $k = |\pi_\sigma^t|$, and so $t = \pi_\sigma^x(k-1) = \pi_t^x(0)$. Starting from α_{k-1} , orient and attach t to form α'_k such that a tile with the same type as $t_{k+1} = \pi_t^x(1)$ can attach to α'_k at a location that is north or west of t . That is, attach t in such a way that the output glue is exposed on the north or west edge of t . Note that this is possible since t is the first tile located outside of $B_c(\mathbf{0})$ and the location of t is to the north of $B_c(\mathbf{0})$. Then, α'_{k+1} is obtained from α'_k by attaching a tile with the same type as t_{k+1} to the north or west of t in α'_k such that the output glue of t_{k-1} is also exposed on the north or west edge. In general, for a step j such that $k+1 < j < n$, α'_j is obtained from α'_{j-1} by attaching a tile with the same type as $\pi_t^x(j-1)$ to the tile t' at the location in $\text{dom } \alpha'_{j-1} \setminus \text{dom } \alpha'_{j-2}$ to the north or west of t' . In other words, we form a path with the same tile types as π_σ^x , only as tiles attach, flip them so that the next tile to attach does so at a tile location that is north or west of the previously attached tile. Denote the resulting path as π' . Figure 7 (left) gives an example of a modified path.

Note that since π_t^x contains at least $4|T|+1$ nodes, π' contains at least $4|T|+1$. By the pigeonhole principle, there must be two nodes \mathbf{v}_0 and \mathbf{v}_1 contained in π' with locations in $B_c(\mathbf{0})$ that correspond to the same tile type in the same orientation. These are depicted as blue tiles in Figure 7. Let $\pi_{\mathbf{v}_0}^{\mathbf{v}_1}$ denote the path in G_α that is the subpath of π_σ^x from \mathbf{v}_0 to \mathbf{v}_1 . Because each tile of π' attaches to the north or west of a previous tile in the path, π' can be modified to extend indefinitely to the north-west since the subpath $\pi_{\mathbf{v}_0}^{\mathbf{v}_1}$ can be repeated an arbitrary number of times. See Table 2(a) for an example of repeating $\pi_{\mathbf{v}_0}^{\mathbf{v}_1}$.

Now, either $\pi_{\mathbf{v}_0}^{\mathbf{v}_1}$ turns or it does not. If it does not turn, then π_σ^x can be modified so that the tile types of the subpath $\pi_{\mathbf{v}_0}^{\mathbf{v}_1}$ are repeated indefinitely




Table 2: Each figure in this table depicts a possible path that can assemble in \mathcal{T} . The yellow tiles make up the seed and the green tiles are a path leading to the repeated tile that allows the path to repeat.

starting from v_0 and this modified path lies on a vertical (or horizontal) line. If each such path π_σ^x from the seed σ to some point x is such that for each pair of distinct nodes of π_σ^x with corresponding tiles of the same type and orientation, the path segment between these nodes does not turn, then every path containing a node in $B_{c'}(\mathbf{0})$ can be modified to assemble an arbitrarily long path that lies on a vertical or horizontal line. Then, it can easily be seen that α is the finite union of semi-doubly periodic sets. This is due to the fact that in the limit, α must consist of σ and finitely many infinitely long paths connected to σ . Each of the infinitely long paths is a semi-doubly periodic set, and there can only be finitely many paths that lie on vertical or horizontal lines that can assemble outside of $B_c(\mathbf{0})$. On the other hand, if there is a π_σ such that the subpath $\pi_{v_0}^{v_1}$ turns, then we can use this turn and the non-orientable nature of tiles in the RTAM to show that α is the finite union of semi-doubly periodic sets.

Let β' denote $\sigma \cup \pi'$. Notice that $\beta' \in \mathcal{A}[\mathcal{T}]$. We now use $\pi_{v_0}^{v_1}$ to form an infinite set of producible assemblies in \mathcal{T} by repeating the tiles of $\pi_{v_0}^{v_1}$ and changing tile orientations by reflecting. Without loss of generality, we assume that output glues of the tiles at locations v_0 and v_1 are on the north edges of the tiles at those locations. (An analogous argument holds when the output glues are on the west edges.) Table 2 depicts this. Table 2 (b) through (h) shows seven of the various paths that can assemble in \mathcal{T} from the path depicted in Table 2(a). We can construct these paths as follows. Since tiles of $\pi_{v_0}^{v_1}$ always attach to the north or west, tile types of $\pi_{v_0}^{v_1}$ can be repeated (in the order that they appear in the path $\pi_{v_0}^{v_1}$ and with the same orientations) indefinitely in the north-west direction to yield a producible assembly. This gives (a) in Table 2. Now, as as tile types of $\pi_{v_0}^{v_1}$ are repeated, at any arbitrary number of repetitions, we can

reflect the tile types of $\pi_{v_0}^{v_1}$ about a vertical axis so that the reflection of the path $\pi_{v_0}^{v_1}$ assembles instead of $\pi_{v_0}^{v_1}$. Let $\hat{\pi}_{v_0}^{v_1}$ denote the reflection of $\pi_{v_0}^{v_1}$ about a vertical axis. Then, tiles of $\hat{\pi}_{v_0}^{v_1}$ can be repeated an indefinite number of times to give (b) through (d) in Table 2. Note that each path of tiles is producible in \mathcal{T} (though perhaps not in the same assembly). Similarly, we can produce paths (e) through (h) in \mathcal{T} by first forming a path where each successive tile attaches to the north-east of the previous tile. This gives an infinite set of paths such that each path belongs to a producible assembly of \mathcal{T} . Denote this set by S . Note that S is an infinite subset of $\mathcal{A}[\mathcal{T}]$.




Fig. 8: A portion of the configuration Γ . This is a map from \mathbb{Z}^2 to T obtained by considering tile types from each path given in Table 2.

Notice that the assemblies in S must be pairwise type-consistent since \mathcal{T} is directed. Consider the map $\Gamma : \mathbb{Z}^2 \rightarrow T$ defined by $\mathbf{x} \mapsto a$ iff $\exists \gamma \in S$ such that $p_T(\gamma(\mathbf{x})) = a$. It should be noted that the fact that assemblies in S are pairwise type-consistent implies that Γ is well-defined. It should also be noted that Γ is equal to the map $p_T \circ \alpha|_{\cup_{\gamma \in S} \text{dom } \gamma}$. Notice that it may be the case that there does not exist an assembly $\phi \in \mathcal{A}[\mathcal{T}]$ such that $\Gamma = p_T \circ \phi$ since orientations of tiles of the same type at a particular location may differ. Figure 8 depicts a portion of this configuration. The points in the domain of Γ consist of the union of a finite number of points and a semi-doubly periodic set. This can be seen by considering the vectors defined by $\mathbf{u} = \mathbf{v}_1 - \mathbf{v}_0 = (a, b)$ and $\mathbf{v} = (-a, b)$.

Using a similar technique that is used to form the assemblies of S , the path $\pi_{v_0}^{v_1}$ can also be repeated (with appropriate reflections) to form type-consistent assemblies containing paths whose domains collectively form a set of points W such that W is the union of four semi-doubly infinite sets. Table 3 shows six examples of such paths. Because the assemblies whose domains make up W are type-consistent, we can once again give a well-defined partial map $p_T \circ \alpha|_W$ from \mathbb{Z}^2 to T . Figure 9 depicts a portion of a configuration given by the map $p_T \circ \alpha|_W$ from \mathbb{Z}^2 to T . Note that any point of \mathbb{Z}^2 is either in W or is in some finite region bounded by points of W .

To finish the proof, it now suffices to show that the set of all points in $\text{dom } \alpha$ (not just the points in W) is the finite union of semi-doubly periodic sets. First, note that except for possibly a finite number of points, any point in $\text{dom } \alpha$ is




Table 3: Each figure in this table depicts a possible path that can assemble in \mathcal{T} . The yellow tiles make up the seed and the green tiles are a path leading to the repeated tile that allows the path to repeat.




Fig. 9: A configuration of tiles that must weakly self-assemble in \mathcal{T} .

contained in some “parallelogram” region R bounded by points of W , where these bounding points are contained in a semi-doubly periodic set with vectors \mathbf{b} , \mathbf{u} , and \mathbf{v} as in Definition 1. Note that for each path whose domain makes up W (see Table 2 for examples of these paths), we can orient the tiles of the path so that the exposed glues of each tile of each path are always oriented the same way. This ensures that if some tile can be placed in the region R , then the same tile can be placed in the region $R_{n,m} = \{\mathbf{r} + n \cdot \mathbf{u} + m \cdot \mathbf{v} \mid \mathbf{r} \in R\}$ for each $n, m \in \mathbb{N}$. See Figure 10 for more details. Therefore, the set of all points in $\text{dom } \alpha$ (even those points that are not contained in W) is the finite union of semi-doubly periodic sets, and hence, α is the finite union of semi-doubly periodic sets.




Fig. 10: (a) Tiles T_1 , T_2 , and T_3 have locations in a region R bounded by points in W . (b) We can choose tile orientations such that tiles with the same types as T_1 , T_2 , and T_3 can bind in regions $R_{n,m}$ for each $n, m \in \mathbb{N}$. (c) A portion of a semi-doubly periodic set containing T_1 , T_2 , and T_3 .

We note that in the proof of Theorem 1, we actually explicitly describe the semi-doubly periodic sets that can assemble in an RTAM system. This is summed up by the following corollary.

Corollary 4. *Let $\mathcal{T} = (T, \sigma, 1)$ be a directed RTAM system. If a set $X \subseteq \mathbb{Z}^2$ weakly self-assembles in \mathcal{T} , then X is a finite union of a box of constant (which depends only on $|T|$ and $|\sigma|$) radius containing the seed and at most 4 (possibly empty) semi-doubly periodic sets. Moreover, each of these 4 semi-doubly periodic sets S_i for $1 \leq i \leq 4$ is described as follows.*

1. S_i is empty.
2. S_i is the finite union of $d \in \mathbb{N}$ horizontal or vertical paths of tiles where d is bounded by $|\sigma|$.
3. S_i is the set obtained from the proof of Theorem 1 by modifying a path that does not lie on a horizontal or vertical line. Note that the number of such sets in this case is bounded by $|T|$.

Proof. Note that the number of paths π as described in the proof of Theorem 1 that may assemble from a single glue exposed by σ is bounded by $|T|$. Therefore,

there is a constant κ that depends only on $|\sigma|$ and $|T|$ such that for the terminal assembly α of \mathcal{T} , $\alpha|_{\overline{B_\kappa(0)}}$ (the portion of α outside of a box of radius κ) is the union of 4 semi-doubly periodic sets. We may take κ to be c' given in the proof of Theorem 1 given in Section 1.

Now we give a corollary that provides evidence to the conclusion that the directed temperature-1 RTAM is not computationally universal. For a tileset T , fix a finite configuration C_T over T , and let $C_T^{\mathbf{v}}$ be the translation of C_T by the vector \mathbf{v} . In other words, C_T is a partial function from \mathbb{Z}^2 to T , and $C_T^{\mathbf{v}}$ is the configuration given by the partial map $C_T^{\mathbf{v}}(x) = C_T(x - \mathbf{v})$ for $x \in \{z \mid z - \mathbf{v} \in \text{dom}(C_T)\}$. Then, let $A(C_T)$ be the set of all seed assemblies σ such that the RTAM system $(T, \sigma, 1)$ is directed and the terminal assembly α_σ of the RTAM system $(T, \sigma, 1)$ contains some translation of the configuration C_T . More formally, $A(C_T)$ is the set of all seed assemblies σ such that the RTAM system $(T, \sigma, 1)$ is directed and the terminal assembly α_σ of RTAM system $(T, \sigma, 1)$ has the property that there exists a vector \mathbf{v} in \mathbb{Z}^2 such that $p_T \circ \alpha_\sigma(x) = C_T^{\mathbf{v}}(x)$ for all $x \in \text{dom}(C_T^{\mathbf{v}})$.

Intuitively, we are considering the set C_T (and its translations) for the following reason. Let $H = \{(i, x) \mid \text{program } i \text{ halts when run on input } x\}$ be the halting set and let M be a Turing machine that outputs a 1 if $(i, x) \in H$. The typical way of expressing the computation of M in tile assembly is as follows. For a fixed tileset T , a seed assembly σ encodes an “input” to the computation, while the “output” of 1 by M corresponds to some configuration being contained in the terminal assembly α_σ of $(T, \sigma, 1)$. For this sense of computation, the following corollary says that the set of seed assemblies that “output” a 1 is a recursive set (not just a recursively enumerable set). This would contradict the fact that the halting set is not recursive. The following corollary is a restatement of Corollary 1.

Corollary 5. *For any tileset T in the RTAM and configuration C_T over T , $A(C_T)$ is a recursive set.*

Proof. (sketch) In each of the cases for the sets S_i described in Corollary 4, to determine if some $C_T^{\mathbf{v}}$ is contained in one of the S_i subassemblies one need only check whether or not a finite portion of S_i contains a translation of C_T , which can be done in a finite number of steps depending only on $|\sigma|$ and $|T|$.

C Proof of Theorem 2

In this section we present the proof that the RTAM is computationally universal at $\tau = 2$. As previously stated, we note that given a Turing machine M , we use Lemma 7 of [2] to obtain a tile set which simulates M using a *zig-zag system*. In fact, as noted in [16], we can find a singly seeded *compact zig-zag system* $\mathcal{T} = (T, \sigma, 2)$ with $\mathcal{A}_\square[\mathcal{T}] = \{\alpha\}$ which simulates M . Then the proof of Theorem 2 relies on showing that any compact zig-zag system in the aTAM at temperature 2 can be converted into a directed RTAM system \mathcal{S} that is “almost” compact zig-zag. The RTAM system that we construct differs from a compact

zig-zag system in that when the length of a row of the growing zig-zag assembly increases by a tile, a strength-2 glue is exposed that allows a tile to bind below the row. This results in the possibility of a single “misplaced” tile per row, but nevertheless, this is enough to simulate a Turing machine. In addition, notice that the orientation of the first tile which binds to the seed tile is such that a strength 2 glue is exposed on either the west or the east. Since the two assemblies obtained from the different binding orientations of this tile are the same (up to reflection), this does not affect the assembly which \mathcal{S} produces. First, we give an example of how to simulate a zig-zag.

Figure 11 shows how we convert a tile set used in a system which simulates a counter in the aTAM at $\tau = 2$ (part (a) of the figure) into a tile set used in a system which simulates a counter in the RTAM at $\tau = 2$ (part (b) of the figure). Part (c) of Figure 11 shows a part of the assembly of the RTAM system \mathcal{S} which simulates the counter.




Fig. 11: Part (a) of this figure shows a tile set which assembles a binary counter in the aTAM, part (b) shows the converted tile set from (a) which assembles a binary counter in the RTAM, and part (c) shows a part of the binary counter assembly from an RTAM system which has the tile set from part (b). A single black box on the side of a tile represents a strength 1 glue, two black boxes on the side of a tile represents a strength 2 glue, and a glue mismatch is represented by a glue label on the side of a tile and the absence of a black box on that side. The tiles with red borders represent “misplaced” tiles.

Proof. (sketch proof of Theorem 2) For the remainder of the proof, we say that two tiles t and t' are of the same form provided that they have glues of the same strengths on the same sides. Given a Turing machine M , we use Lemma 7 of [2] to obtain a tile set which simulates M using a zig-zag system. Furthermore, as noted in [16], we can find a singly seeded compact zig-zag system $\mathcal{T} = (T, \sigma, 2)$ with $\mathcal{A}_{\square}[\mathcal{T}] = \{\alpha\}$ which simulates M . We assume that \mathcal{T} grows to the north. Let $\mathcal{S} = (S, \sigma', 2)$ be an RTAM system defined as follows.

First, for clarity, we rewrite the glues contained on tiles of T in terms of their complementary glues. We use the convention that if a glue g appears on

the north or east side of a tile, then we simply write g . If a glue g appears on the south or west of a tile, then we rewrite that glue as g' . Next, we set σ' to be the same as the tile which composes σ . It follows from the constructive proof of Lemma 7 in [2], that for any $t \in T$ such that t has a strength 2 glue on its south, it must be of the same form as tile A shown in Figure 12(a) up to reflection. For each $t \in T$ with a strength 2 glue on the south, we create a tile $t' \in S$ which consists of the same south glue as the one on t , but has a copy of the east/west glue of t such that t' has identical east and west glues. The form of t' is that of tile A' shown in Figure 12(b). Henceforth, we refer to tiles of the same form (up to reflection across the vertical axis) as the tile labeled B in Figure 12(a) as *corner tiles*. For each corner tile $t \in T$, we form a tile $t' \in S$ which consists of the same east and west glues as t , but has a copy of the north/south glue of t such that t' has identical south and north glues. Thus, in our RTAM tile set S , corner tiles in T take on the form of the tile labeled B' in Figure 12(b). Notice that tiles which require cooperation to bind are necessarily oriented upon binding. Consequently, for all tiles $t \in T$ which bind cooperatively (i.e. those that are not of the form A or B shown in Figure 12(a)), we add tile t to our RTAM tile set S .




Fig. 12: (a) Tiles of this form (upto reflection) allow a row of zig-zag growth in the aTAM to increase in length by one tile before proceeding to the growth of the next row. (b) Tiles of this form (upto reflection) allow a row of “almost” zig-zag growth in the RTAM to increase in length by one tile before proceeding to the growth of the next row.

Now, we show that the northern most row of tiles in the assembly obtained from \mathcal{S} contains the final configuration of the tape of M and its final state just as it is in \mathcal{T} . Because tiles which bind cooperatively in \mathcal{S} always orient themselves in the same manner in which they appear in \mathcal{T} , we only examine the cases where τ strength glues are used in binding. Up to reflection, all of the cases of tiles which bind with strength 2 are of the same form as tiles $A' - C'$ in Figure 12(b). Tiles that are of the same form as the tile labeled A' in Figure 12(b) do not pose a problem since the extra glue it has relative to its counterpart in t is strength 1 and exposed on the east/west edges of the assembly in such a way that no other tiles can bind to expose a glue with which it can cooperate.

Next, we examine the binding of a tile of the same form as B' in Figure 12(b). Upon binding, a tile of the same form as B' will allow for the binding of a tile which its counterpart in T does not. That is, it will allow for the binding of a tile to its south. But, the constructive proof of Lemma 7 in [2] is such that only tiles of the same form as A' in Figure 12(b) can flip and bind to the south of

the B' tile. Now, observe that due to the zig-zag nature of growth in S , the tile located diagonally to the south-east of the misplaced A' tile is of the same form as the tiles B' or D' in Figure 12(b). Notice that this means the misplaced A' tile cannot cooperatively place a tile with any other glue.

It follows from the compact zig-zag Turing machine simulation construction used in [16], that tiles of the form C' appear only once in α . More specifically, tiles of the same form as C' always attach to the seed tile. Thus, in our case it does not matter which orientation C' has upon binding since the two assemblies obtained from allowing C' to bind in its two orientations are the same up to reflection.

Consequently, the “misplaced” tiles in the assembly produced by \mathcal{S} , that is the tiles that differ from the tiles in the assembly produced by \mathcal{T} in location or label, are tiles that bind to the south of corner tiles. Since this does not interfere with the simulation of M , \mathcal{S} is computationally universal. To prove the second statement of the theorem, notice that the constructed RTAM system is directed.

D Proof of Theorem 3

In this section, we give the full proof of Theorem 3.

Proof. We prove Theorem 3 by contradiction. Therefore, assume that for some $n \in \mathbb{Z}^+$ such that $(n \bmod 2) = 0$, there exists an RTAM system $\mathcal{T} = (T, \sigma, 1)$ such that $|\sigma| = 1$ and \mathcal{T} weakly self-assembles an $n \times n$ square S . It must therefore be the case that for some subset of tile types $B \subseteq T$, for all $\alpha \in \mathcal{A}_\square[\mathcal{T}]$ there exist $r \in R$ and $\mathbf{v} \in \mathbb{Z}^2$ such that for $\alpha_r = F(\alpha, r, \mathbf{v})$, $\alpha_r^{-1}(B) = S$ (that is, every terminal assembly contains an $n \times n$ square of “black” tiles, and no “black” tiles anywhere else).




Fig. 13: The shortest path between two opposite corners of an $n \times n$ square has length $2n - 1$.

We choose an arbitrary $\alpha \in \mathcal{A}_{\square}[\mathcal{T}]$ and refer to the four corners of S in α as \mathbf{a} , \mathbf{b} , \mathbf{c} , and \mathbf{d} , with \mathbf{a} being the southwest and moving clockwise to name the remaining corners. (See Figure 13 for an example using an 8×8 square.) Since α must be a stable assembly, there must exist a set of (possibly overlapping) paths P in the binding graph of α which contain the tiles at all 4 corners. Let π_a^c be a shortest path which contains both corners \mathbf{a} and \mathbf{c} (there may be more than one), and note that $|\pi_a^c| \geq 2n - 1$ because the tiles are located on a grid. (Two possible such paths can be seen in Figure 13 as the red and yellow paths.) Let the uppercase letters A , B , C , and D represent the quadrants which contain the corners with the corresponding lowercase labels. It must be the case that π_a^c contains at least one tile in either quadrant B or D , else it could not reach quadrant C . This is due to the facts that n is even and that π_a^c is a path through a grid graph where diagonal travel is not allowed. Without loss of generality, assume that π_a^c enters quadrant B .

We now know that some path π_a^c through the binding graph of α connects \mathbf{a} and \mathbf{c} and also contains at least one point in quadrant B . Additionally, we know that within the overall binding graph of α , π_a^c must somehow also be connected to the tile at \mathbf{d} . Let π_d^{ac} be a path in the binding graph which contains both \mathbf{d} and some tile in π_a^c . Note that this may just be the tile at \mathbf{d} since π_a^c may contain that location. We now define path π' as the longest of the following two paths formed by possibly combining subpaths of π_a^c and π_d^{ac} : 1) a path which connects \mathbf{d} and \mathbf{c} and contains a tile in B , or 2) a path which connects \mathbf{a} and \mathbf{d} and contains a tile in quadrant B . For ease of discussion, we will talk about π_a^c as being a directed path from \mathbf{a} to \mathbf{c} , while noting that we don't know the actual ordering of its growth. Such a π' must exist because the point at which π_d^{ac} intersects π_a^c must be either 1) before it has entered B , which yields case 1 for π' , 2) after it has left B , which yields case 2 for π' , or 3) within B in which case either holds. Note that π_a^c could perhaps enter and leave B multiple times, which would only strengthen our argument, but for simplicity we will simply make use of the first time it enters quadrant B and the last time it leaves quadrant B for the above argument.

Without loss of generality, we will assume case 1 for path π' , and note that we now have found the following path in the binding graph of α . Path π' includes the tile at location \mathbf{d} , contains at least one tile in quadrant B , and also includes the tile at location \mathbf{c} . (See Figure 14 for an example.) For a square of dimension n , it must be the case that path π' contains a minimum of n pairs of tiles bound to each other on their east/west edges.

Now we check to see if the seed σ is contained within π' . If not, we define π'_σ as the minimum path in α 's binding graph which contains σ and some tile in π' , else π'_σ is simply σ . Now we define a new valid assembly sequence, α' for \mathcal{T} as follows.

Let \mathbf{x} denote the location of the intersection of π'_σ and π' , and π_x^c be the portion π' from \mathbf{x} to \mathbf{c} and let π_x^d be the portion of π' from \mathbf{x} to \mathbf{d} . Without loss of generality, we will assume that the tile on π'_σ which binds to σ attaches to the north of σ . (If this is not the case, then all following directions can simply be rotated by the same angle as the difference of the direction of the first binding




Fig. 14: Example paths in an 8×8 square (i.e. of even dimension). White tiles represent tile types not in B (and thus tiles not in S), black tiles represent path π'_σ , dark grey tiles represent path π_d^{ac} , and π_a^c is represented by a combination of tiles of all colors. Path π' connects d and c by passing through x (which is in quadrant B) and is outlined in red.




Fig. 15: Every possible pairing of an input and output side of a tile on a path. Those in grey have as input sides either south or west, and as output sides either north or east. Those in yellow do not, but are labeled with the minimal reflection necessary to orient them so that they do.

from north.) Starting from σ , α' attaches one tile at a time from π'_σ so that output glues are exposed on the north or east edge of each tile. In other words, as tiles attach, flip them so that the next tile to attach does so at a tile location that is north or east of the previously attached tile. (Figure 15 shows how this must be possible for each tile.) This causes all inputs along π'_σ to be on the south or west sides of tiles, and outputs on the north and east. Furthermore, place the tile from location x so that the exposed glue used for the binding of π_x^c is on its north or west, and the exposed glue used for the binding of π_x^d is on its south or east. Figure 16 shows how this can be done. (Note that depending on where and from which direction π'_σ intersects π' , it could be the case that the output sides which bind to π_x^c and π_x^d are instead south or east, and north or west, respectively. However, again this doesn't change the argument as the following directions can be rotated appropriately, so we assume the output for π_x^c is on the north or west for ease of discussion.)

Now we assemble the path π_x^c as follows. Starting from the tile of location x , attach one tile at a time so that the output glues are exposed on the north or west edge of each tile. This results in a version of π_x^c which grows strictly up and to the left. Such growth is possible because of the set of reflections possible and the fact that no tile previously placed from π_x^c or π'_σ can block the placement. The fact that blocking can't occur is due to the fact that the stretched version of π_x^c grows strictly up and/or to the left so all previous tiles placed for π_x^c cannot block, and after the first tile placed after the final tile of π'_σ , the newly forming




Fig. 16: Every possible pairing of a south or west input side and 2 output sides of the final tile of π'_σ . Those in grey have one output side on the north or west, and the other on the south or east. Those in yellow do not, but are labeled with the minimal reflection necessary to orient them so that they do.

path is either completely above π'_σ or to the left of the topmost tile of that path. Next we assemble the path π_x^d in an analogous manner, but down and to the right. Starting from the tile of location \mathbf{x} , attach one tile at a time so that the output glues are exposed on the south or east edge of each tile. For the same but reflected reasons as for the growth of the modified version of π_x^c , π_x^d can be grown in this way. See Figure 17 for an example of these “stretched” paths.




Fig. 17: Example paths π_σ^x , π_x^c , and π_x^d stretched out when assembled by α .

The tile types which are at the ends of π_x^c and π_x^d were at locations \mathbf{c} and \mathbf{d} in α , so they must be of types which are in $B \subseteq T$ (i.e. “black” tile types). However, since there are $\geq n$ pairs of tiles connected via east/west glues in π' and it is now maximally stretched, it must be $\geq n + 1$ tiles wide (i.e. from leftmost to rightmost tiles). Thus, this is a producible assembly which has “black” tiles at a distance which is further apart than any two points in the square S . Therefore

α' does not weakly (or, therefore, strictly) self-assemble the $n \times n$ square S and so neither does \mathcal{T} . This is a contradiction, and thus \mathcal{T} does not exist.

E Proof of Theorem 4

In this section, we give the details of the proof of Theorem 4.

Proof. The general scheme for the construction can be seen in Figure 18. Essentially, the seed, which is in the center of the square, presents some glue on both the north and south, and presents some second glue on both the east and west. Above the seed, a series of $(n - 1)/2$ hard-coded tile types form a vertical column, with the east and west edges of all tiles exposing the same glue. Using the same tile types, a reflected version of the upper column grows downward to the bottom of the square. To both the east and west of that column, a hard-coded series of another $(n - 1)/2$ tile types form reflected rows out to the sides of the square. This requires $1 + 2((n - 1)/2) = n$ tile types and strictly self-assembles an $n \times n$ square.

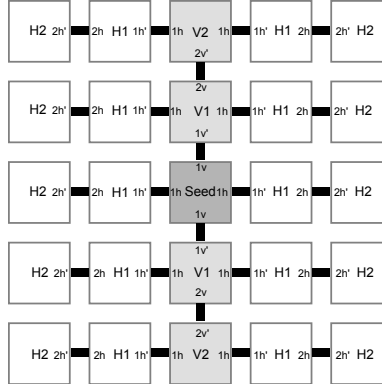


Fig. 18: A 5×5 square built by an RTAM TAS at $\tau = 1$ using only 5 tile types.

To prove Theorem 4, we provide the following construction which demonstrates how to build an RTAM TAS which strictly self-assembles an $n \times n$ square for any odd n . The basic idea is that we place the seed tile in the middle of the square, and from both its top and bottom it grows copies of a reflected column of tiles which grow to the top and bottom of the square, respectively. This requires one tile type for the seed and $(n - 1)/2$ tile types for the column (which grows in both directions). Each of the seed and the tiles of the vertical column have the same glue on both east and west sides. From each of those a row of $(n - 1)/2$ unique tile types grows to the left or right boundary of the square. This construction is robust to any valid rotation of the tiles and strictly self-assembles an $n \times n$ square using exactly $1 + 2(n - 1)/2 = n$ tile types.

In Figure 19 we show the templates used to generate the necessary tile types to self-assemble an $n \times n$ square for odd n , and in Figure 18 we show an example




Fig. 19: The tile type templates for building a square whose sides are odd length in the RTAM at $\tau = 1$ using n unique tile types.

of the terminal assembly for $n = 3$. For all odd n , the seed is exactly as in Figure 19. Then for each m where $0 < m < (n - 1)/2$, we create a unique tile type from the tile type template labeled “Vx”. We do this by replacing each “x” (i.e. those of the tile label and the south glue) with the value “ m ”, and the “y” of the north glue with the value “ $m + 1$ ”. (Note the “prime” markings of some glues which denote they are complementary to the “unprimed” versions, and for all tile types created from these templates we keep those markings unchanged, which is important in restricting the potential interactions.) In this way, we create the tile types for the central column between the seed and the top or bottom most locations. To make the top/bottom tile type, we start with the “Vz” template and replace the “z” of the label and south glue with the value $(n - 1)/2$. To make the tile types for the horizontal rows which grow outward from the central column, for each m where $0 < m < (n - 1)/2$ we create a unique tile type from the tile type template labeled “Hx”. We do this by replacing each “x” (i.e. those of the tile label and the west glue) with the value “ m ”, and the “y” of the east glue with the value “ $m + 1$ ”. In this way we create the tile types for the horizontal row between the central column and left or right most locations. To make the left/right tile type, we start with the “Hz” template and replace the “z” of the label and west glue with the value $(n - 1)/2$. This completes the creation of exactly n tile types: 1 for the seed, $2(((n - 1)/2) - 1)$ for the interiors of the column and rows, and 1 each for the top/bottom tile type and left/right tile type. The RTAM TAS \mathcal{T} consisting of this tile set, a seed consisting of a copy of the seed tile in any valid orientation at the origin, and $\tau = 1$ strictly self-assembles an $n \times n$ square and is directed. This is because the only tiles which can attach to the north and south of the seed are properly reflected versions of the “V1” tile type, and then to the north/south of those two tiles properly reflected copies of “V2”, etc. until copies of the “V $(n - 1/2)$ ” tile type cap off the north and south growth of the column. Regardless of the reflections of the “Vx” and “Vz” tiles across the vertical axis, the only subassemblies that can grow to the left and right are the arms consisting of the “Hx” and “Hz” tile types. All

producible terminal assemblies of \mathcal{T} are equivalent and in the shape of an $n \times n$ square.

F Proof of Theorem 5

In this section, we give the full proof of Theorem 5.

Proof. We prove Theorem 5 by contradiction. Therefore, assume that for some $n \in \mathbb{Z}^+$ such that $(n \bmod 2) = 1$, there exists an RTAM system $\mathcal{T} = (T, \sigma, 1)$ such that $|\sigma| = 1$ and $|T| < n$, and \mathcal{T} weakly self-assembles an $n \times n$ square S . It must therefore be the case that for some subset of tile types $B \subseteq T$, for all $\alpha \in \mathcal{A}_\square[\mathcal{T}]$ there exist $r \in R$ and $\mathbf{v} \in \mathbb{Z}^2$ such that for $\alpha_r = F(\alpha, r, \mathbf{v})$, $\alpha_r^{-1}(B) = S$.

We choose an arbitrary $\alpha \in \mathcal{A}_\square[\mathcal{T}]$ and, using the notation of Figure 13, note that in the binding graph of α , there is a path π_a^c which connects the tiles in locations \mathbf{a} and \mathbf{c} (i.e. opposite corners of the square). Furthermore, the tile types of the tiles at locations \mathbf{a} and \mathbf{c} must be in $B \subseteq T$, and tiles of types in B can be no further apart from each other (in the plane) than a Manhattan distance of $2n - 1$. Since the seed σ may not be on π_a^c , define the location \mathbf{x} to be either 1) the location of σ if it is on π_a^c , or 2) the location of the intersection of π_a^c and the shortest path in the binding graph which contains σ and some tile in π_a^c . Note that such a location \mathbf{x} must exist since α must be connected. Define the path π_σ^x as the path from σ to \mathbf{x} and note that if σ is on π_a^c then π_σ^x consists of a single tile. Also, we can now define π_x^a and π_x^c as the paths from \mathbf{x} to \mathbf{a} and \mathbf{c} , respectively. (See Figure 20 for a possible example of such paths in α .)




Fig. 20: Example paths π_σ^x (darkest grey) and π_a^c (a combination of light grey, dark grey, and white tiles) in α which weakly self-assembles a 9×9 square (i.e. of odd dimension). White tiles represent tile types not in $B \subseteq T$, while grey and dark grey tiles represent “black” tile types in B .

We will now define a valid assembly sequence α' in \mathcal{T} which builds an assembly α' which consists of modified versions of the paths π_σ^x and π_a^c , and which places tiles of types in $B \subseteq T$ (i.e. “black” tiles) at points further apart than $2n - 1$, proving that \mathcal{T} does not weakly self-assemble S . Let $k = |\pi_\sigma^x|$ and let

α' start from σ and first build a modified version of π_σ^x in the following way. If $k = 1$, stop. Else, without loss of generality assume that $\alpha(\pi_\sigma^x(1))$ attaches to the north of σ in α . (If it binds on the east, south, or west, for the following argument used to construct α' , rotate all directions used clockwise by 90, 180, or 270 degrees respectively.) For each i where $0 < i < k$, place a tile of the same type as at $\alpha(\pi_\sigma^x(i))$ but reflected so that its input side is either south or west, and its output side is either north or east. Note that for any pair of input and output sides, it is possible to reflect a tile to meet this criteria (see Figure 15). Place the tile of that type and orientation so that it binds with the $(i - 1)$ th tile of the newly created path. Note that this also must be possible because, following from σ , every tile is oriented so that the side used as its output side in π_σ^x is north or east and has a glue which matches that of the newly placed tile on its (now) south or west side, respectively, and since the newly forming path grows only up and/or to the right, no previously placed tile can block it from binding. The result of this sequence of placing the tiles from π_σ^x is a “stretched” version of the path π_σ^x which grows only up and/or right from the seed σ , grown by the valid assembly sequence α' .

We now extend α' to grow a stretched version of π_a^c by growing stretched versions of π_x^a and π_x^c . We will build one of π_x^a or π_x^c by stretching it up and to the left, and the other by stretching it down and to the right. (See Figure 21 for an example.) Since the final tile of π_σ^x is included in π_a^c , it must be the case that it has two output sides (to bind to each of π_x^a and π_x^c). Also, recall that its input side must have been either south or west. We modify α' so that the final tile of π_σ^x is oriented with one output side on either the north or west, and the other on either the south or east. See Figure 16 for all possible orientations of this tile. Without loss of generality, assume that in this orientation, the side which serves as input to π_x^a is on either the north or west, and that serving as input to π_x^c is on either the south or east. (If it is the opposite, the directions used to grow the stretched π_x^a and π_x^c can simply be swapped and the argument will still hold.) We will now build a stretched version of π_x^a up and to the left, and π_x^c down and to the right. To do this, we extend α' to attach the tiles of π_x^a , moving along that path as it appears in α and placing the same tiles in the same order, but always orienting them so that each has as its input side either its south or east, and as its output side either its north or west, allowing each successive tile to correctly bind to its predecessor on the path. As before, this is possible because of the set of reflections possible and the fact that no tile previously placed from π_x^a or π_σ^x can block the placement. The fact that blocking can't occur is due to the fact that the stretched version of π_x^a grows strictly up and/or to the left so all previous tiles placed for π_x^a cannot block, and after the first tile placed after the final tile of π_σ^x , the newly forming path is either completely above π_σ^x or to the left of the topmost tile of that path. The stretched version of π_x^c is grown analogously, with all input sides being on the north or west and output sides to the south or east.

Note that in the special case where $x = a$ or $x = c$ in α , meaning that either the seed is in one of those corners, or the path from the seed to the path π_a^c

At this point, α' grows assembly α' which is a path that looks like that in Figure 21 and which starts from the seed and grows a stretched version of π_σ^x , then grows a stretched version of π_a^c (as connected stretched versions of π_x^a and π_x^c) from the end of that path. Since in α , π_a^c connects the bottom-left corner of the square S with the top-right corner, it must be the case that $|\pi_a^c| \geq (2n-1)$. Since $|T| \leq (n-1)$, and $2(n-1) = 2n-2$, by the pigeonhole principle we know that along π_a^c (and therefore also its stretched version), it must be the case that at least 3 tiles of some same type appear.

35




Fig. 22: One possible configuration of a portion of the stretched paths π_σ^x (black) and π_x^a and π_x^c (grey). Three reflected copies of the same tile type appear on π_x^a and π_x^c (lightest grey). A path using different input sides can be formed between the top and bottom copies of the repeated tile type, entering from the a and b sides, respectively.

We now modify α' so that it builds the stretched copies of π_σ^x , π_x^a , and π_x^c until it is about to place the tile at f . At this point, it rotates the tile for that position so that the side which serves as input for the tile at g is facing either south or east, and places that tile. Then, since the tile at g was above and/or to the left of that at f , the side it used as output is now oriented so that it can be used as output from the copy at g . We add to α' the tile placements which now make a copy of the path π_g^f , again orienting the final tile to allow yet another copy. We make α' add $2n$ copies of π_g^f . Then, we let the final tile on the final copy of π_g^f be oriented so that the final portion of the original stretched copy of π_x^c can grow. At this point, α' , which is still a valid assembly sequence in \mathcal{T} , forms the stretched copy of π_σ^x , as well as the stretched copy of π_x^c , with the addition that some interior portion of π_a^c has been “pumped” so that the overall distance between the tiles at the end of the path are at a Manhattan distance at least $2n$ from each other. This is because the stretched path π_a^c is at least that long, containing $2n$ copies of π_g^f , and it grows so that each tile placed makes the endpoints strictly further from each other. Further, the tiles placed at the endpoints of the stretched version of π_a^c are of the same types as those placed at a and c in α . Since those were at the corners of the square S , they were tile types from $B \subseteq T$ and must therefore be inside the square S . However, in S , no two points are at a Manhattan distance greater than $2n - 1$ from each other. Thus, α' does not weakly (or therefore strictly) self-assemble S and neither does \mathcal{T} . This is a contradiction, and thus Theorem 5 is proven.

G Proofs of Theorems 6, 7, and 8

G.1 Proof of Theorem 6

Consider the shape S depicted in Figure 23a. By Corollary 3, there is a singly seeded RTAM system that weakly self-assembles S . Therefore, to complete the proof, we show that this shape cannot be strictly self-assembled by a singly seeded system in the RTAM by contradiction. Suppose that $\mathcal{T} = (T, \sigma, \tau)$ is a singly seeded system that strictly assembles S , and let α in $\mathcal{A}_{\square}[\mathcal{T}]$ be a terminal assembly such that $\text{dom}(\alpha) = S$. Without loss of generality, assume that the seed for \mathcal{T} is at a location shown in red in Figure 23a. Then, by applying a single tile reflection, we can modify the assembly sequence of α to obtain a producible assembly β depicted in Figure 23b, which contradicts the fact that S strictly assembles in \mathcal{T} .




Fig. 23: (a) A shape S that cannot strictly self-assemble in the RTAM. (b) A producible assembly that can assemble in any system which can also produce an assembly whose domain is S (upto translation and reflection).

G.2 Proof of Theorem 7




Fig. 24: (a) A shape S that cannot be self-assembled in the RTAM by a directed system. (b) A producible assembly that can assemble in any system which can also produce an assembly whose domain is S (upto translation and reflection).

Consider the shape S depicted in Figure 24a. First, to see that S can be strictly self-assembled in the RTAM by a singly seeded system, consider the tile set shown in Figure 25.




Fig. 25: A tile set T such that a system $\mathcal{T} = (T, \sigma, \tau)$, with σ a single tile assembly consisting of the tile labeled S , strictly self-assembles the shape S shown in Figure 24a. Note that all of the glues depicted in this figure have strength τ , and that assembly of the shape S requires the reflection of these tiles.

Then, to complete the proof, we show by contradiction that S cannot be strictly self-assembled by a directed singly seeded system in the RTAM. Suppose that $\mathcal{T} = (T, \sigma, \tau)$ is a directed singly seeded system that strictly assembles S , and let α in $\mathcal{A}_{\square}[\mathcal{T}]$ be a terminal assembly such that $\text{dom}(\alpha) = S$. Without loss of generality, assume that the location of the seed in α is at a location shown in red in Figure 24a.

Let t_i denote the tile of α located at positions L_i . Note that the south glue of either t_1 or t_5 is bound to α with strength τ . Without loss of generality, assume that the south glue of t_1 is bound to α with strength τ . If t_1 is bound to t_2 with strength τ , t_2 is bound to t_3 with strength τ , and t_3 is bound to t_4 with strength τ , then by modifying the assembly sequence for α and applying the appropriate reflections, we can give an assembly that is producible in \mathcal{T} , and is a total of 11 tiles wide. An example of such an assembly is shown in Figure 24b. This contradicts the assumption that \mathcal{T} strictly self assembles S . On the other hand, if the strength of the bond between t_i and t_{i+1} is less than τ for $1 \leq i \leq 4$, then the south glues of t_1 and t_5 must each be bound with strength τ . Therefore, the assembly α' depicted as gray tiles in Figure 26 must be producible in \mathcal{T} . In the




Fig. 26: A shape S that cannot be self-assembled in the RTAM by a directed system. Strength- τ glues are shown with two bars, the rest of the glues can be either strength-1 or strength-2 as long as the assembly is τ -stable.

following argument, without loss of generality, we assume that for every glue, g , on a tile in T , there is a matching glue g' on some (possibly the same) tile in T . Then, since two tiles must be able to attach so that they occupy the two tile locations to the west of t_1 , after applying a reflection of t_1 , these same tiles, call them a and b , must be able to attach with strength τ to the east to t_1 with the easternmost tile, b , located at L_3 . Similarly, (prior to a and b binding) three tiles, call them a' , b' , and c' , must be able to attach with strength τ to the west of t_5 , with a' located at L_4 , b' located at L_3 , and c' located at L_2 . Notice that b' must have two strength- τ glues. Hence, nondeterministically, either b or b' will be at location L_3 . Finally, note that b and b' cannot have the same tile type.

This follows by contradiction; suppose that b and b' are of the same type. Then, after binding, b must expose a strength- τ glue. However, when the westernmost tile that attaches to the west of t_1 is b , this strength- τ glue would allow for a tile to bind, giving a producible assembly with a domain that is not contained in S . This contradicts the assumption that \mathcal{T} strictly self-assembles S . Finally, since in any terminal assembly of \mathcal{T} , either b or b' will be at location L_3 , and b and b' do not have the same tile type, we conclude that \mathcal{T} is not directed.

G.3 Proof of Theorem 8

Let S denote the shape shown in Figure 24a. Then, the proof follows from the proof of Theorem 24. Note that in proof of Theorem 24, if b is placed at L_3 , then the west edge of a' and the east edge of b must have different glues. Similarly, if b' is placed at L_3 , then the east edge of a and the west edge of b' must have different glues. Therefore, any singly seeded RTAM system that strictly self-assembles S is not mismatch-free.

H Proof of Theorem 9

First, we show that if S is ϵ -symmetric, then there is a singly seeded mismatch-free RTAM system that strictly assembles S . Let T_S be a tree for S which satisfies Definition 3. We use T_S to define a tile set T and give a seed tile σ such that the system $\mathcal{T} = (T, \sigma, 1)$ strictly self-assembles S , and then we argue that for any terminal assembly α in \mathcal{T} , the domain of α is equal to S (up to translation and reflection). T and σ are obtained as follows.

First, notice that we can give a temperature-1 aTAM system $\mathcal{T}' = (T', \sigma', 1)$ that assembles the shape S with the properties (1) \mathcal{T}' is a singly seeded directed system, (2) for $\alpha' \in \mathcal{A}[\mathcal{T}']$, the binding graph of α' is exactly the tree T_S , and (3) when the seed tile σ' is the single tile located at some point v , then for each point p of S , there is a unique tile t' in T' such that $\alpha(p) = t'$. More intuitively, when the seed tile of \mathcal{T}' is placed at v , then assembly proceeds by the binding of unique tiles equipped with glues specific to each point of S .

Now, we define the tiles of T to be copies of the tiles in T' and define σ' to be σ . Notice that in the RTAM, if we assume that no tiles are reflected before binding as assembly proceeds, then the terminal assembly, α , that results has a domain equal to S . Furthermore, since the binding graph of α is T_S , a reflected tile, t , of α lies on some axis, a_t say, and corresponds to some vertex v_t of that axis. Then, since T_S is ϵ -symmetric and appropriately reflected tiles of T can still bind to t even when t is reflected prior to binding, the domain of a terminal assembly of \mathcal{T} will either (1) be equal to S (This is the case that the axial branches beginning from v_t are equivalent upto reflection about the axis a_t . See Figure 27b.), or (2) be equal to a reflection of S about the line corresponding to the axis a_0 (This is the special case that the axial branches beginning from v_t are not equivalent after a reflection about a_t . See Figure 27c.). In either case, it follows that S strictly assembles in \mathcal{T} .




Fig. 27: An example of the assembly of a portion of an ϵ -symmetric shape which assembles in \mathcal{T} . Relative to Figure (a), Figures (b) and (c) give two different producible assemblies of \mathcal{T} .

Now we show that if there is a singly seeded mismatch-free RTAM system that strictly assembles S , then S is ϵ -symmetric. Let $\mathcal{T} = (T, \sigma, 1)$ be a singly seeded mismatch-free RTAM system that strictly self-assembles S , and let $\alpha \in \mathcal{A}_{\square}[\mathcal{T}]$ be some terminal assembly of \mathcal{T} .

First, we note that the binding graph of α must be a tree, which we denote by T_{α} . This follows from the fact that if the binding graph contains a cycle, then there are infinitely many producible assemblies in \mathcal{T} . See Figure 28 for an example. Moreover, T_{α} must be ϵ -symmetric. This can be shown by contradiction as follows.




Fig. 28: (a) An example of a binding graph that contains a cycle. (b) An example of a binding graph of an assembly where tiles with type of those tiles belonging to the cycle shown in (a) are repeated an arbitrary number of times using appropriately reflected tiles.

Suppose that T_{α} is not ϵ -symmetric. Then, there exists an axis a of T_{α} such that T_{α} is not off-by-one symmetric across a . In other words, there are at least two vertices, which we denote by v_1 and v_2 , of T_{α} which lie on a such that the

branches of a which begin from v_1 are not symmetric across a and the branches of a which begin from v_2 are not symmetric across a . Now, we can modify the assembly sequence of α by reflecting the tile, t , located at v_1 prior to binding, allowing reflected tiles to assemble reflections of the subassemblies which are bound to t , and keeping the original assembly sequence otherwise. Now, as tiles of the reflected subassemblies originating from t bind, either a mismatch occurs or the reflected subassemblies completely assemble. If a mismatch occurs, we arrive at a contradiction since \mathcal{T} is mismatch-free. If the reflected subassemblies completely assemble, then we note that no translation and reflection of $\text{dom}(\beta)$ is equal to S . (This is because we essentially chose a sequence where the two asymmetric branches across a adopted reflections across a such that the sides of each which would be on the same side of a in S are now on different sides of a .) This contradicts the assumption that \mathcal{T} strictly assembles S . Therefore, T_α is ϵ -symmetric. Hence S is ϵ -symmetric.




Fig. 29: An example of possible binding graphs of a shape that is not ϵ -symmetric. (a) and (b) depict binding graphs of two distinct assemblies.

I Proof of Theorem 10

Let S be a finite connected shape in \mathbb{Z}^2 and S^2 be S at scale factor 2. Note that we can give a temperature-1 aTAM system $\mathcal{T}' = (T', \sigma', 1)$ that assembles the shape S with the properties (1) \mathcal{T}' is a singly seeded directed system, (2) for α' in $\mathcal{A}_\square[\mathcal{T}']$, the binding graph of α' is a tree $T_{\alpha'}$, (3) the location of the seed tile σ' corresponds to a leaf of $T_{\alpha'}$, and (4) for each point p of S , there is a unique tile t' in T' such that $\alpha(p) = t'$. For an example of such an assembly α' see Figure 30.

We now give an RTAM system \mathcal{T} based on \mathcal{T}' such that a terminal assembly α of \mathcal{T} can be obtained from α' by replacing single tiles of α' by 2×2 blocks of tiles, and thus assembles S^2 . aTAM tiles and corresponding RTAM tiles that give rise to this block replacement scheme are described in Figure 31.

As 2×2 blocks assemble, cooperation is used to ensure that the orientations of the tiles of a 2×2 block are fixed relative to the orientation of the seed.




Fig. 30: An example terminal assembly giving a shape assembled by an aTAM system \mathcal{T}' .

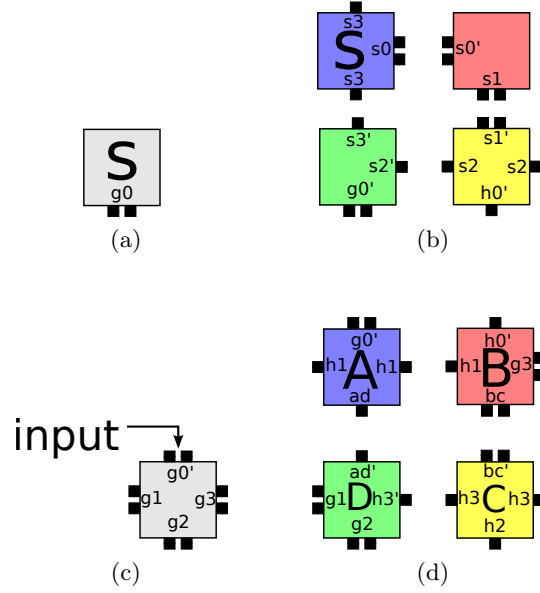


Fig. 31: (a) An example seed tile of the aTAM system \mathcal{T}' . (b) Tiles that represent the seed tile of (a) used to define the RTAM system \mathcal{T} . The seed of \mathcal{T} consists of only the tile labeled S . (c) An example tile of the aTAM system \mathcal{T}' that attaches via the glue $g0'$. Note that it may be the case that not all glues g_1 , g_2 , and g_3 shown here are exposed. (d) Tiles that represent the tile of (c) used to define the RTAM system \mathcal{T} . Notice that tiles bind in the order A , B , C and D . The tile shown in (c) has “output” glues g_1 , g_2 , and g_3 exposed. In the case where one or more of these glues is not exposed by a tile, we obtain tiles analogous to those in (d) that represent this single tile as follows. The glues g_1 , g_2 , and g_3 of the tiles of (d) are exposed if and only if the respective glues of the tile in (c) are exposed. For example, if the tile in (c) exposes a g_1 glue, then so will the tile D of (d). Otherwise, D will not expose a g_1 glue. Note that the cooperative binding of B fixes the orientation of B . Similarly, the cooperative binding of D fixes the orientation of the tile D .

Figure 32 depicts a terminal assembly of the RTAM system \mathcal{T} based on the aTAM system \mathcal{T}' whose terminal assembly is shown in Figure 30.




Fig. 32: An example terminal assembly giving a shape assembled by an RTAM system \mathcal{T} based on the aTAM system \mathcal{T}' whose terminal assembly is shown in Figure 30.